

Toward a complete predictive rendering pipeline for numerical prototyping and cultural heritage applications

Romain Pacanowski

▶ **To cite this version:**

Romain Pacanowski. Toward a complete predictive rendering pipeline for numerical prototyping and cultural heritage applications. Informatique [cs]. Université de Bordeaux (UB), 2023. ⟨tel-04380582⟩

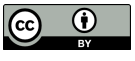
HAL Id: tel-04380582

https://inria.hal.science/tel-04380582v1

Submitted on 8 Jan 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons CC BY 4.0 - Attribution - International License

Toward a complete predictive rendering pipeline for numerical prototyping and cultural heritage applications

By

ROMAIN PACANOWSKI



Ecole Doctorale de Mathématiques et d'Informatique
UNIVERSITÉ DE BORDEAUX

A dissertation submitted to the University of Bordeaux

MARCH 30, 2023

Dedication and acknowledgements

To Zelik Pacanowski

Table of Contents

Page

List of Tables **vii**

List of Figures **ix**

1 Introduction **1**

1.1 Contributions on Qualified Measurements . . . 2
1.2 Appearance Modelisation and Edition . . . 3
1.3 Light Transport Simulation . . . 5
1.4 Outline of the manuscript . . . 5

I Image-based Acquisition Systems **7**

2 *La Coupole*: A SV-BRDF measurement device for non-planar objects **9**

2.1 Designing *La Coupole* . . . 10
2.2 La Coupole Acquisition System . . . 12
2.3 A short Review of Appearance Acquisition Systems . . . 15
2.4 Measurement Procedure Overview . . . 16
2.5 First Reconstruction Results . . . 17
2.6 Calibration . . . 18
2.6.1 Overview . . . 18
2.6.2 Accuracy and Repeatability Assessment . . . 21
2.7 Discussion . . . 22
2.7.1 Normal Direction Uncertainty . . . 22
2.7.2 Acquisition Time and Storage Constraints . . . 22
2.8 Future Work . . . 23
2.8.1 Improving *La Coupole* . . . 23
2.8.2 Efficient Representations . . . 25

II Reflectance Models **27**

3 Rational Approach for BRDF Measurements Approximation **31**
3.1 Rational BRDF Definition . . . 31
3.2 Fitting Measurements Data with Rational BRDF . . . 32
3.3 Results on some Matusik Materials . . . 34
3.3.1 Error Analysis in BRDF space . . . 35
3.3.2 Error Analysis in Image space . . . 36
3.4 Informal Comparisons with Recent Approaches . . . 38
3.5 Future Work . . . 40

4 Combining Microfacet and Diffraction theories **43**
4.1 Introduction . . . 43
4.2 Overview . . . 46
4.3 A Two-Scale BRDF Model . . . 46
4.3.1 A diffractive Microfacet . . . 47
4.3.2 The Microfacet lobe of the two-scale model . . . 50
4.3.3 The Diffraction lobe of the two-scale model . . . 51
4.4 Overall Fitting Results . . . 54
4.5 Results for our conductor model . . . 55
4.6 Multi-Layer Extensions . . . 56
4.7 Fitting Technique . . . 58
4.8 Comparisons with the He et al. model . . . 59
4.9 Discussion and Future Work . . . 59

Appendices **63**
4.A The polarization factor Q . . . 63
4.B Precomputed Data . . . 63

III Efficient Predictive Rendering **65**
4.3 Monte-Carlo Methods for light transport simulation . . . 68

5 Per-pixel Optimized Balancing for Multiple Importance Sampling **73**
5.1 Motivation . . . 73
5.2 Second-order Variance Reduction . . . 75
5.3 Theoretical analysis . . . 77
5.4 Convergence speed for α computation . . . 77
5.5 α clamping . . . 77
5.6 Results . . . 78

5.7 Discussion . . . 79
5.8 An improved version . . . 81
5.8.1 Three Strategies Balancing . . . 82
5.8.2 Malia Implementation Overview . . . 83
5.9 Short-Term Future Work . . . 83

IV Future Directions **89**

6 Future Directions **91**
6.1 Main Objectives . . . 91
6.2 Axis 1: Appearance measurement with quantified uncertainties . . . 93
6.2.1 To validate the predictions of models . . . 93
6.2.2 For Cultural Heritage Preservation . . . 95
6.3 Axis 2: Light transport simulation and appearance synthesis . . . 95
6.3.1 Radiometric simulation . . . 96
6.3.2 Predictive Rendering . . . 96
6.4 Axis 3: Inverse models for appearance prediction . . . 97

List of external references **103**

Bibliography **121**

List of Tables

Table **Page**

5.6.1 Efficiency, Variance and Lab error comparisons between different variance reduction techniques 80

List of Figures

Figure **Page**

1.1 Istropic BRDF imaging setup . . . 2
1.2 Diffraction detection . . . 3
1.3 Retro-reflective materials samples and measurements . . . 3
1.4 Appearance of layered materials with correlated nano-particles . . . 4
1.5 Efficient Sampling of 4D Luminaire . . . 5

2.1 Examples of degradations on fabric material . . . 10
2.2 Belt of the Chinese general . . . 11
2.3 La Coupole . . . 12
2.4 Technical drawings of La Coupole . . . 13
2.5 Dibond Support . . . 13
2.6 Dome II acquisition device . . . 16
2.7 Overview of the reconstruction results . . . 17
2.8 Reconstruction of the prophylactic shirt . . . 18
2.9 Calibration Targets . . . 19
2.10 Overview of the camera extrinsic parameters calibration . . . 20
2.11 Automatic SV-BRDF Segmentation . . . 25

3.1 Number of coefficients and maximum degrees found for each rational fit . . . 34
3.2 Lens-flare-like acquisition artefacts . . . 36
3.3 Specular materials' BRDF data . . . 37
3.4 Maximum relative error of the rational approximation . . . 38
3.5 Rendering Comparisons for Rational BRDFs . . . 39

4.1 Complete Fitting Results from Ngan et al. . . . 44
4.2 The failure of Gaussian-shaped lobe . . . 44
4.3 Butler BRDF Measurements . . . 45
4.4 Optical Path Difference on a rough height-field surface . . . 47
4.5 Diffraction effects depend on projected vector $\boldsymbol{f}$ . . . 49
4.6 The exponential power distribution (EPD) is controlled by two parameters β and p. . . . 51

4.7 Approximation of the energy renormalization term A . . . 53
4.8 Evaluation of the convolution $(S * D)(\theta_h)$ using s and g functions . . . 54
4.9 RMSE on the different tested BRDF models for all materials of the MERL database. 54
4.10 Lobe shapes Visualization for `blue-metallic-paint2` . . . 55
4.11 Image space comparisons between our model and Löw model . . . 56
4.12 Illustration of the lobe width variation wrt. to the light wavelength . . . 57
4.13 Plastic and Subsurface variants of our diffraction model . . . 57
4.14 Example of a material where our model performs well . . . 58
4.15 Example of a material where the RMSE for our model is larger . . . 58
4.16 Rendering Comparisons for different diffraction BRDF Models . . . 60
4.2.1 GPU vs. CPU compute capabilities comparisons . . . 67
4.2.2 Denoising in real-time rendering . . . 68
4.2.3 CNN-based denoiser comparisons . . . 69

5.5.1 Number of samples M to estimate α per scene . . . 78
5.5.2 Effect of clamping α for each scene. . . . 78
5.6.1 Illustration of per-pixel α computations for balancing . . . 79
5.6.2 Test Scenes for Variance Reduction Technique . . . 84
5.6.3 Visualization of per-pixel variance. . . . 85
5.8.1 Renderings with Global illumination using precomputed α-maps . . . 86
5.8.2 Comparison between one iteration and 8 using global illumination to compute α . . . 87
5.8.3 Efficiency comparison in the presence of strong indirect lighting . . . 88
5.8.4 Efficiency comparisons when using 3 sampling strategies . . . 88

6.1.1 Examples of ancient materials . . . 91
6.1.2 Industrial Materials . . . 92
6.1.3 Natural Materials . . . 92
6.1.4 Radiometric Nomenclature for multi-layered systems . . . 92
6.1.5 Overview of the proposed methodology . . . 93
6.2.1 A simulated Speedshape with nano-coating compared to a real one . . . 94
6.4.1 1D, 2D, and 3D Photonic Crystals . . . 98
6.4.2 Chirped and Bragg Systems . . . 99
6.4.3 Appearance of a Bragg System . . . 99
6.4.4 Periodic systems with elliptical nano-particles . . . 100
6.4.5 Industrial and Biological examples of multi-layered materials . . . 101

CHAPTER 1

INTRODUCTION

Although synthetic images surrounding us (special effects, video games, advertisements, etc.), often appear to be as plausible or even realistic, they rarely present physical guarantees and cannot be used for predictive purposes which is necessary in the architectural field or in the production of manufactured objects. My research activity focuses on the synthesis of predictive images and is organized around the three main steps of the opto-numerical chain that allow to generate such images:

1. **Acquisition and measurements** of the different quantities related to appearance (shape, light and matter) in order to ensure that the models developed are consistent with the reality of physical phenomena;
2. **Modeling and representation** of the quantities involved in the rendering algorithms;
3. Improvement of the algorithms and techniques **simulating the transport of light**.

During my master research and my PhD, I have been interested in the last two stages of the opto-numerical chain. I started working on the first stage (acquisition) during my post-doctorate at CEA-CESTA and also when I was recruited at the *Laboratoire de Photonique, Numérique et Nanosciences* (LP2N). Since my recruitment at INRIA (December 2020) my research activity continues to span all stages of the opto-numerical chain.

As explained through out this manuscript, my research direction is moving towards a combination of metrological measurements of the appearance with predictive image synthesis methods (also known as predictive rendering). For example, diffraction effects may originate from the materials or from the sensor; the metrological approach then allows to answer with more certainty about the origin of such visual effects. This metrological approach is essential for inverse design.

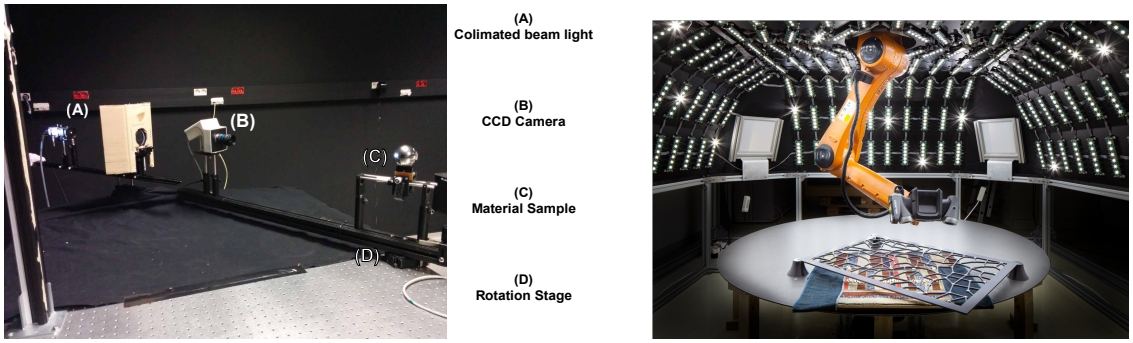


Figure 1.1: **Left:** Photograph of the isotropic BRDF measurement setup. **Right:** Photograph of *La Coupole* mounted with a 3D Scanner during the shape acquisition step.

Finally, the inverse design of a material from its desired optical properties and final appearance is one of my main research goals for the next 10 years.

The next three short sections (1.1, 1.2, 1.3) give an high-level overview of my research activities (2005-2020), for each of my research axes (Acquisition, Modeling, Light Transport Simulation). Finally, Section 1.4 presents the organisation of the rest of the document.

1.1 Contributions on Qualified Measurements

In collaboration with Ramon Hégédus, we published a metrological study [1], both analytical and numerical, on imaging devices that measure the Bidirectional Reflection Distribution Function (BRDF) when the material support is spherical (see Figure 1.1). The originality of our 3D BRDF instrument is not so much in its design or its realization with more powerful hardware than previous work, but in the rigorous study of the uncertainty associated with each element of our device, contrary to the original experiment of Matusik et al. [MPBM03]. One of the major problems with this type of device is that it is necessary to perform high dynamic photography (HDR) in order to properly capture the reflection properties. However, the HDR process also generates aberrations related to the diffraction of light by the diaphragm of the lens, thus making certain parts of the image unusable. This observation led us to develop a new technique [2, 3] to detect (see Figure 1.2) diffraction effects in an HDR image from a confidence threshold set by the experimenter. Unlike previous approaches based on de-convolution (e.g., [Ric72]), we obtain images with a lower maximum error and uncertainty per pixel; this makes our method particularly suitable in the context of metrology. This preliminary results (cf. [4]) allowed us to conceive, in the framework the *Materials* ANR Project (2015-2020), an opto-numerical acquisition system (cf. Figure 1.1) in order to digitize the appearance of old clothes stored in the reserves of the Museum of Ethnography of Bordeaux. The opto-numerical acquisition system, later named *La Coupole*, is designed to acquire the geometry as well as the appearance of objects of any shape and of large sizes (about $2m^2$ and up to 70 cm high). Unlike previous work (e.g., [SSW+14]), the originality of *La Coupole* (cf. Chapter 2) is to acquire separately the shape (geometry) and reflection properties of large objects. The performance of *La Coupole* in terms of spatial and

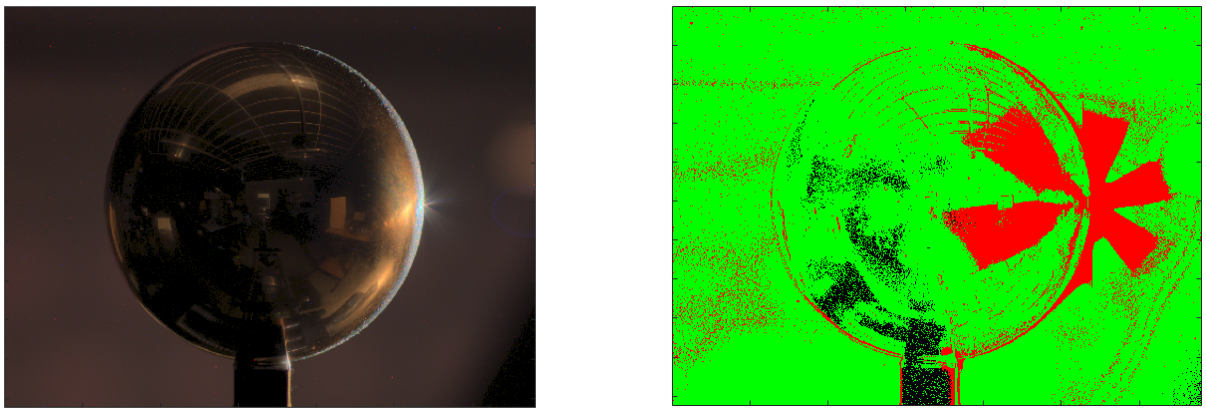


Figure 1.2: **Left**: HDR image with diffraction effects, especially in the form of stars. **Right**: Result of our method [3] that classifies pixels in red if they are affected by diffraction effects and in green otherwise.

angular resolutions and acquisition time exceeds existing instruments. Indeed, it has a resolution of 50 microns (equivalent to the device of Filip et al. [FVH+13]) and is therefore more efficient than the existing ones [HP03, MMS+04, NKRS15, SSWK13a, SWR+11] except for [SSK03]. On the other hand, the measurement density allowed by Sattler et al. [SSK03] setup is two orders of magnitude lower than the metrologically qualified device that we propose.

1.2 Appearance Modelisation and Edition

My work on modelling and appearance editing has focused a lot on the reflectance of materials defined by the BRDF. During my PhD. I was interested in direct editing of reflections by a user using drawings ([5]), in editing the appearance by modifying the incident illumination via perturbations of its frequency distribution [6] or its amplitude ([7]). More recently, in the framework of Boris Raymond's thesis, we were interested in editing anisotropic reflections [8]

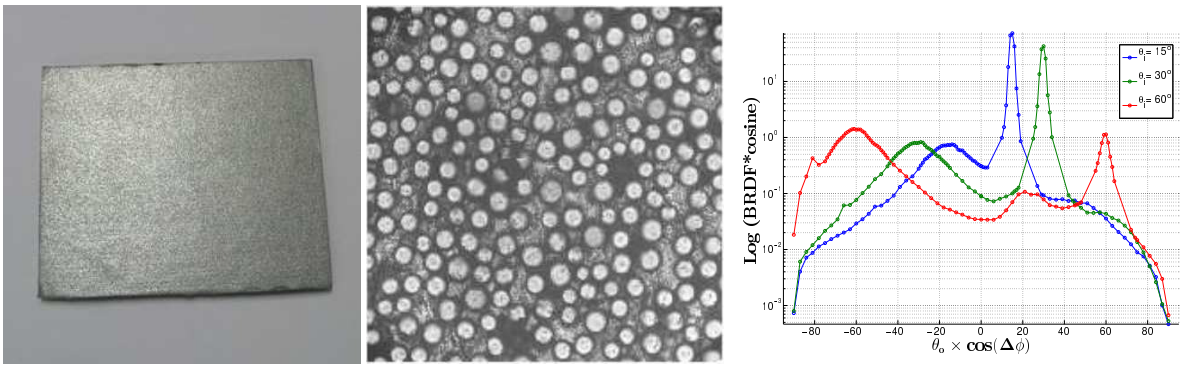


Figure 1.3: **Left**: Retro-reflective material example. **Center**: Optical Surface Profiler (NewView 7300) images revealing the micro-balls structure of the retro-reflective mater. **Right**: BRDF corrected by the cosine factor measurements in the plane of incidence ($\Delta\phi = 0$ or $\Delta\phi = 180$) where the retroreflective zone corresponds to positive abscissa values (i.e., $\theta_0 \cos\Delta\phi \geqslant 0$).

on a surface. During my post-doctorate, we proposed a new data representation [9] in order to better approximate analytically the BRDF measurements and their importance function to facilitate their sampling. More recently, we have proposed a new BRDF model (see Chapter 4) that effectively takes into account diffraction effects, which are too often ignored.

From BRDF measurements made during my post-doc at CEA-CESTAS, with Laurent Belcour [10] and colleagues, we introduced an extension to the microfacet model to take into account retro-reflection (cf. Figure 1.3) effects. Our main contribution was to introduce a new vector, the back vector that is the equivalent of the half-vector for retro-reflection effects.

From 2015 to 2020, with Kévin Vynck, Pascal Barla, Arthur Dufay and others, I worked on the direct and inverse modeling of the appearance of nano-structured materials in order to map the different effects that can be obtained at the macroscopic scale while ensuring that the material is manufacturable. The first step [11] of our work was to develop a BRDF model (cf. Figure 1.4) that predicts quantitatively the optical response of such layered-materials with nano-particles.

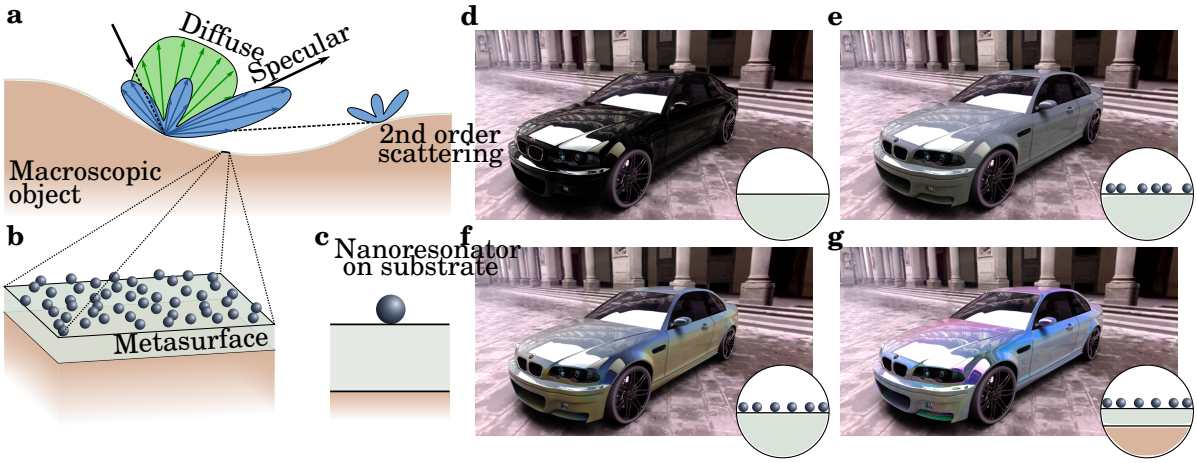


Figure 1.4: **a-c**, Multiscale nature of the visual appearance of a multi-layered materials with nano-particles, from the individual particle resonantly interacting with the substrate interfaces **(c)** to the ensemble of particles interacting together to form collective scattering effects **(b)**, to the macroscopic surface that can experience multiple internal reflections **(a)**. **d-g**, Synthetic images, obtained with our new BRDF model, of a macroscopic object in a realistic lighting environment. In **d**, the car body is made of an raw tinted glass, leading to a purely specular and dominantly black appearance. In **e**, the same surface is covered by Ag particles of radius $r = 90$nm distributed at random positions. Scattering by the particles reduces specular reflection and yields a diffuse, nearly isotropic, grey appearance. In **f**, a control of structural correlations in the particle positions leads to the emergence of a striking new appearance. In **g**, a structured substrate (SiO2 layer of thickness $h = 400$nm on a Si substrate) is used, considerably enriching the observed diffuse colours.



Figure 1.5: **Left:** Light pattern projected on a wall from a car headlight that can be modeled as a 4D (two dimensions for spatial variations and two for angular variations) light source, a light field acquired through many images. **Center and Right:** Our sampling technique [12] allows to sample very efficiently this type of light field to achieve interactive (7 fps on a workstation with an Intel Core i7 920 with 6 GB and a GTX 580 with 1.5GB) framerate rendering. In this case, the light field was measured using the Lumigraph parametrization with 11×9 images (each image has a 256^2 pixels). The center and right images are obtained with the same captured data thanks to our technique that allows to modify interactively the position and the orientation of the captured light field.

1.3 Light Transport Simulation

During my master and my PhD, I was interested in a volumetric representation of indirect incident illumination [13] in order to accelerate the generation of realistic images. During Heqi Lu's PhD thesis [Lu14], we worked on efficient real-time sampling of time-varying light sources [14] as well as 4D light sources (cf. Figure 1.5) by proposing an appropriate sampling technique [12] or to optimize [15] the combination (Multiple Importance Sampling) of sampling techniques (material vs. light source). With Xavier Granier and David Murray we further improved [16] M.I.S., thanks to some precomputations (more details in Chapter 5) In the framework of Arthur Dufay's CIFRE thesis (Technicolor 2014-2017), we improved the pre-visualization of realistic images, in the context of the creation of special effects, by proposing a new random sample generation algorithm via a micro-jittering technique [17] allowing a significant gain in performance while being adapted to graphics hardware architectures.

1.4 Outline of the manuscript

The rest of the manuscript is organized in a non-conventional way. Each chapter describes a contribution that is detailed and discussed inversely proportionally to to the size of the associated publications. For example, the contribution on Multiple Importance Sampling (cf. Chapter 5) is quite long due to the fact that only a short paper at Eurographics was published. I chose not to discuss some recent publications (e.g., [18, 19]), which contributions are very well described in the Ph.D. manuscript (cf. [Bat21]), and of course, too old ones. I also chose to not detail some technical contributions (e.g., the Importance Sampling of the combined microfacet-diffaction BRDF model or the micro-jittering approach [17]).

Unsurprisingly, the chapters are grouped into three parts, one for each stage of the opto-numerical chain. Furthermore, I conclude each chapter with some short term future work related to its content. Finally, Chapter 6 presents my research directions in the medium and long terms.

Part I

Image-based Acquisition Systems

CHAPTER 2

La Coupole: A SV-BRDF MEASUREMENT DEVICE FOR NON-PLANAR OBJECTS

In July 2015, we obtained a funding from the French National Research Agency (ANR) for our project *Materials*. Combining the expertise from scientists coming from Computer Science (Maverick and Manao INRIA teams), Optics (LP2N), Printing (Océ Print [1]) and Ethnography (MEB), the goal of the project was to scan cultural artefacts conserved by the Musée d'Ethnographie de Bordeaux (MEB) and to print them in 2.5D using newly developed technologies from Océ Print.

The overall idea of the project started by an initial demand from our ethnography colleagues who were (and are still) interested in knowing if digital acquisition could help them preserve and exhibit their fabric artifacts. Indeed, the whole collection of fabric materials, stored in the basement of the museum, were rarely shown to the public due to the different degradations suffered through time.

The origin of the degradations (cf. Figure 2.1) are diverse:

- inadequate storage units (high temperature and hydrometry promoting the proliferation or moths) they were conserved in,
- the transportation from their former museum (Musée du Trocadéro in Paris) to their current location (MEB),
- necessary manipulations by humanities and social sciences experts and researchers.

Inside the Materials project, the LP2N was naturally in charge of the design, development and construction of an acquisition device that could capture the geometry and the reflectance of

[1] Océ Print, a subsidiary from Canon, was in charge of the fabrication through printing technologies. Unfortunately, Océ Print went into bankruptcy in June 2019.

the different artifacts. As shown in Figure 2.2, the artifacts are cloths (or pieces of cloth) coming from Asia (Laos, Vietnam, China) and hand-manufactured at the end of the 19th century. Cloths were chosen as a case study since these artifacts are among the most challenging ones.

The following sections describe and present the design and realization of our solution named *La Coupole* (formerly known as *Le Dôme*).

2.1 Designing *La Coupole*

The design of our acquisition system was guided by four types of constraints :

- Physical due to the space requirement
- Financial and time
- Preservation
- Virtual reproduction (rendering quality).

Physical Constraints.

An initial request from the museum's curators was to avoid the transport of the cloths to the LP2N laboratory. Instead we had at our disposal an "acquisition" room next to the storage room where our device should be installed.

The dimensions ($4 \times 4\,\mathrm{m}^2$) of the acquisition room, especially the height under the ceiling (2.40m) combined with the width of the staircase to access it, are the main external physical constraints that impose the overall wingspan of *La Coupole*. Naturally, a second constraint is the size of the cloths we had to acquire. Even if a cloth can lie on a support, it is still a large object (its diameter is around 1m) requiring a large support.

Preservation Constraints.

Acquiring the shape and reflectance properties should be possible under the following constraints from the museum curators:

- No chemical products should be put on the cloths.



Figure 2.1: Examples of degradations on fabric material. **Left:** degradation due to insects. **Middle:** beat bark **Right:** degradations due to misuse.



Figure 2.2: **Left:** Photograph of the belt of the (**Center**) Chinese general is one the pieces of cloth we acquired during our measurements campaign in summer 2019. The belt length is approximately 60 cm for a 30 cm width. **Right:** Close-up photograph of another acquired cloth, the "wedding vest". Chinese General Photograph Credit Dr. Matignon around 1900.

- No UV or infrared lights should be used.
- No mechanical (twisting, cutting, etc.) deformation should be applied.

These constraints led us to choose optical technologies thar are inherently non invasive for the material sample. However, even if light wavelengths from the visible domain are allowed, a certain amount of light quantity should not be exceeded. To that end, we followed the recommendations [NB11] from the Centre de Conservation du Quebec which specify that the quantity of light should not exceed 50 lux during 20 hours [2] of lighting per year. In practice this means that, once the appearance is acquired, the artifact should be put in the store for one year.

Financial and Time Constraints.

Obviously our initial budget from the ANR was limited but it was tripled thanks to funding provided by a regional grant driven by Inria Bordeaux Sud-Ouest and a *Contrat Plan Etat Région* (CPER) grant driven by the Optics Institute (IOGS). Overall, our financial hard constraint to build our device was 150 000 euros and the device needed to be completed by July 2019 in order to be ready for the exhibition organized and planned by the MEB for Fall 2019.

Reproduction Constraints.

Reproduction constraints are also a quantitative objective regarding the quality of the synthetic

[2] This is an overall limitation. One could also use 100 lux during 10 hours or any lux × time equivalent.



Figure 2.3: La Coupole, two measurement configurations. **Left:** The 3D scanner is mounted on the robotic arm to acquire the shape (geometry) of the object. **Right:** Instead of the 3D scanner, the high-speed camera is mounted on the robotic arm to acquire the hemispherical reflectivity with four white LED panels or the SV-BRDF with 1080 LEDs.

image one would like to achieve from the data acquired. In agreement with our partners, we decided that our reproduction goal should be to mimic at least the resolution and precision of the human eye. According to Grenon [Gre09], the human vision system of a healthy human being can accommodate at 25cm with a resolution of $100\mu m$. Moreover, the angular resolution of the human eye is driven by the diameter of the iris that ranges from 2mm (myosis) to 8mm (mydriasis). We compromised with a 4mm iris diameter that defines a 1 degree[3] angular resolution. The equivalent in photography is a 50mm lens with a f-number ($f\#$) equal to 11.

2.2 La Coupole Acquisition System

All the afore mentioned contraints led us to build *La Coupole* (cf. Figure 2.3) that consists mainly in the following parts:

- A 6-axis robot from Kuka
- A 3D scanner from Creaform
- A high-speed camera from Ximea
- 1080 LEDs that act as directional light sources
- Four flat white LED panels (from Airis company) that act as a uniform hemispherical source
- A support, mounted on a rotation stage, on which the artifact will lie on
- Mechanical rails forming a gallows that supports the robot
- 5 aluminium panels that support the LEDs. The front panel is mobile and serves as an access door.

[3] Or more precisely: $2 \cdot \arctan\left(\frac{0.2}{25}\right) = 0.92$ degree.

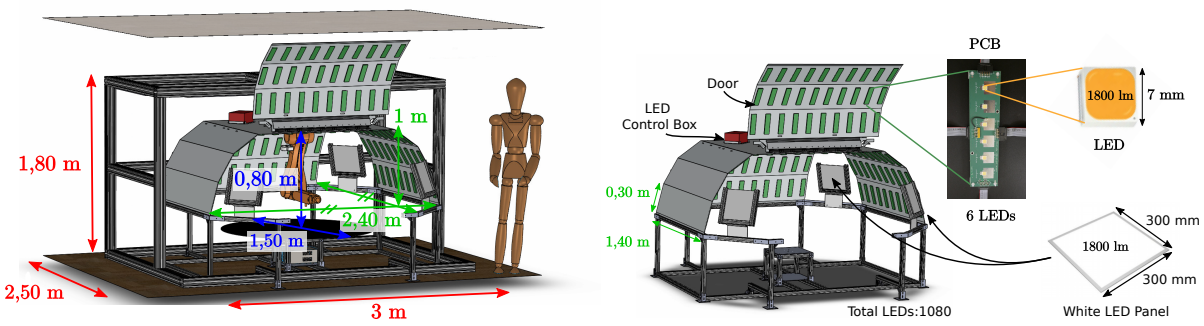


Figure 2.4: **Left:**. Dimensions of *La Coupole*. **Right:** Focus on the two types of light sources. The 4 white LED panels are used to generate an hemispherical lighting whereas the single LED mounted on PCBs are used for directional lighting. All lumens (luminous flux) numbers are the ones provided by the manufacturers.

The total cost of *La Coupole* is approximately 120 000 euros. Its weight is slightly inferior to 1000 kg (see also Figure 2.4).

The **support** (cf. Figure 2.5) for the material sample consists of three parts. From top to bottom, first, a flat disk which diameter is 1.48m, made of Dibond[4] that supports the artifact, second, a stainless disk (8mm thick), which is hollowed, as much as possible, to reduce its weight (35 kg) is attached to the last part, the rotation stage. The Dibond has the advantage to be lightweight whereas the stainless disk increases the rigidity of the whole system.

The **rotation stage** from Thorlabs (HDR-50M model) supports up to 50 kg with a bidirectional accuracy of 820μrad and a bidirectional repeatability of 350μrad. The rotation stage is used for the LED calibration as well as for virtually increasing the number of LEDs.

The 6-axis robot is used to move either the camera or the 3D scanner. With a maximum reach of 1.1 meter the robot (KUKA Model Agilus kr10r1100sixx C) can reach all positions of the Dibond board. According to the manufacturer, its pose repeatability is 30μm.

[4]Dibond is actually a brand. A Dibond consists of a sheet of two aluminium layers (0.3mm) that encapsulate a polyethylene or a mineral core.

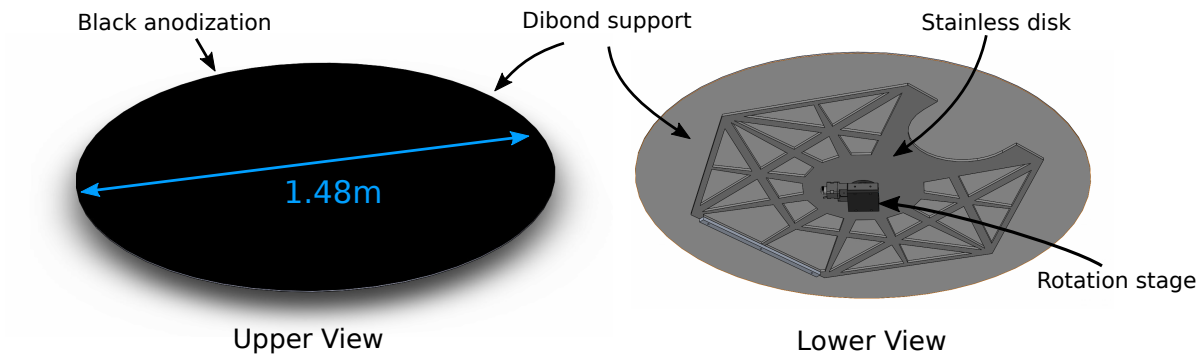


Figure 2.5: The support for the material sample consists of three parts: **Left to Right:** A black anodized Dibond support, a stainless reinforcement disk and a rotation stage.

The 3D laser Scanner (Model HandyScan 700 from Creaform) is used to acquire the shape (geometry) of the sample. Its resolution is announced at 50μm (distance between two points) with an accuracy of 30μm. In another words, the position of point measured from the scanner is known up to 30μm. The speed of the scanner (480 000 points per second) makes is very competitive compared to photogrammetry. Indeed, a 1m^2 surface can be acquired in 15 minutes. Furthermore, our preliminary tests in 2016 showed that photogrammetry techniques were having big difficulties with the acquisition of cloths. From a metrology point of view, the HandyScan precision can be traced back to the meter standard delivered by the *Institut National de Métrologie*. The Handy scanner needs to see retro-reflective stickers to orientate itself. These stickers are pasted on a calibration target positioned around or above the material sample (cf. Figure 2.9).

The lighting system consists of 1080 LEDs and 4 flat white LED panels (cf. Figure 2.4). The four LED Panels (manufactured by Airis) generate an almost perfect uniform lighting and serve different purposes. During the calibration step, they are used to calibrate the spatial sensitivity of the sensor (see below) but they are also used during the measurement step to acquire the hemispherical reflectivity of the sample.

The 1080 LEDs are distributed on 180 electronic Printed Circuit Board (PCB). The 180 PCBs are fixed on aluminium structure that surrounds the Dibond board. Each LED is a 7mm square, which luminosity is announced to be around 1800 lumens. Each LED can be turn on during 1ms (electrical current rise constraint) and up to 2 seconds (heat dissipation constraint[5]). An electronic control box (see Figure 2.4) allows to control each LED and also triggers the camera. This hardware triggering mechanism permits to obtain good synchronization performance.

The camera is used to acquired the optical properties of the material sample. Our choice (Ximea CB120CG-CM) was guided by the reduction of the acquisition time. Therefore, we would like an acquisition system that can sustain the lighting rate of the LEDs, which ranges from 10ms to one second. Hence, the camera should at least support 100 frames per second. The Ximea RGB-CMOS sensor has a resolution of approximately 4200×3000 pixels where each pixel size is 5.5μm and the signal to noise ratio is quite good with 42 dB. The high speed rate of the camera is achieved through a specialized x4 PCI-express card, which bandwidth is 20 gigabytes per second. The chosen **lens** is a 35 mm Sigma HSM Art. The minimal focus distance is stated to be 30 cm. With a lens aperture set to at least f#8, we can achieve the one degree angular resolution required to mimic the human eye. The maximum aperture authorized by the lens is f#16. The combination of the Sigma lens and the Ximea sensor is an imaging system that acquires a 14.5×19.3cm^2 rectangular surface where each image pixel, at normal incidence, represents a square of 47μm on a side.

[5]At two seconds, the LED lifetime starts to drop quickly due to the produced heat.

2.3 A short Review of Appearance Acquisition Systems

Before describing the measurement and calibration procedures of our device we review in this section some previously built appearance acquisition systems. We voluntarily focus on spatially-varying BRDF. An overview for more complex radiometric quantity (e.g., BSSRDF) can be found in the state of the art report of Fridsvad et al. [FJM+20].

BRDF measurement devices (more or less automated) have been developed since the sixties. For example, the less-famous paper from Torrance and Sparrow [TS66] already describes the measurement of a radiometric quantity called the "Biangular Reflectances", which is actually a 2D BRDF as introduced later by Nicodemus in its nomenclature [NRH+77]. Historically speaking, their paper reports 3D BRDF measurements made by Fritz Taler [Tha03] in 1903 using petrol lamp as light source, or even older measurements in 1888 from Messerschmitt [Mes88].

Early BRDF measurement devices are called spectro-goniophotometers. They measure the BRDF for one single light and view directions at a time. The light source (a laser for example) is collimated on the material sample and a photodiode measures the signal. The support of the sample can be a goniometric-cradle and additional rotation stage are used to cover the 4 dimensions of the BRDF. With the development of digital photography, measurement devices for the BRDF have become imaging systems (e.g., [Mar98, MPBM03]). Finally, a recent and complete overview of BRDF measuring instruments is available in Weinmann and Kleiny [WK15] as well as in Guarnera et al. [GGG+16].

It is worth noticing that without imaging technologies it would be too time consuming to measure a spatially-varying BRDF densely.

Mc Allister [McA02] and Sattler et al. [SSK03] were among the first to measure a SV-BRDF **on a planar surface** using a camera. Their systems were quite slow due to the many mechanical movements that are needed to orientate the material sample. To overcome this limitation, recent systems (e.g., [HHN+17]) are doing acquisition in parallel (aka multiplexing) by adding several cameras and light sources.

However, a major limitation remains for appearance acquisition. Objects (e.g., museum artifacts) are rarely planar. Consequently, more recent systems (e.g., [HLZ10, SWRK11, SSWK13b, KNRS13]) combine structured lighting to acquire the shape of the object with multiple light sources and/or cameras.

The solutions introduced by Schwartz (cf. Figure 2.6) and by Köhler are related to ours since they rely on a dome to place multiple lights and cameras. However, compare to our solution, the resolution of the view direction is inferior to ours since the number of cameras is fixed (e.g., 11 for [SSWK13b]). The geometry is acquired during a first step using structured lighting projected by video Cloths were chosen as a case study since these artifacts are among the most challenging onesprojectors. In summary, previous dome-based systems are limited in terms of camera resolution and object size.

The object size limitation comes mainly from the choice of the field-of-view. It is a compromise

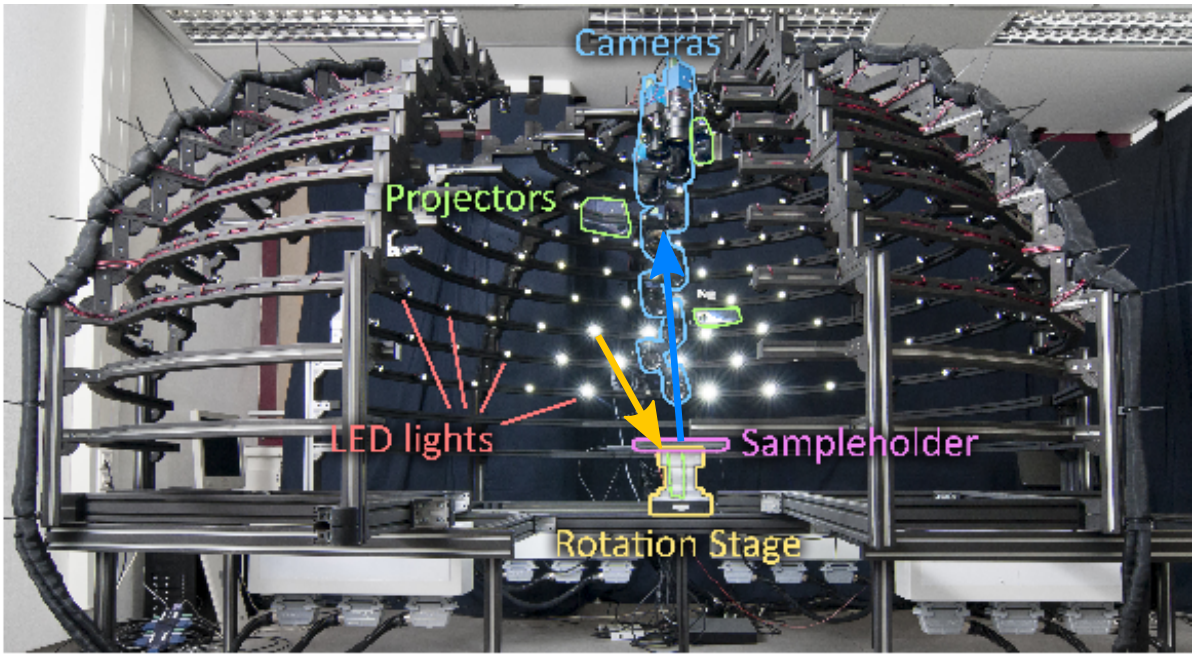


Figure 2.6: The Dome II from [SSWK13b] is built with 11 cameras, 188 LEDs and 4 video projectors.

between the spatial resolution of the measurement and the object size. In practical terms, for *La Coupole* imaging 1m^2 at a resolution of $100\mu\text{m}$ would require a sensor of 100 million pixels per view direction and adequate lens to have sufficient depth-of-field yielding in a too costly camera.

Finally, it is worth mentioning that the main difficulty when acquiring the shape and the reflectance is to re-project the acquired photos on the shape geometry. Our solution to this problem is presented in Section 2.6 where we give more insights on the calibration procedure. The system proposed by [HLZ10] is quite original: two robotic arms are moving around the object and each arm carries a camera and video projector, which is used either to project a structured pattern or to illuminate the object. The main disadvantages of the system is its relative slowness and the maximum size of the object is also limited by the camera field-of-view.

2.4 Measurement Procedure Overview

Once the calibration is done (cf. Section 2.6) the measurement is performed in three steps:

- Acquisition of the object shape thanks to the 3D scanner,

- Acquisition of the hemispherical reflectance accomplished by replacing the 3D scanner with the Ximea Camera and by switching on the white panels. One or multiple exposures (HDR) image(s) is acquired per camera position.



Figure 2.7: Reconstruction results obtained for 8 cloths. Additional videos are available online.

- Acquisition of the spatially-varying reflectance. This is accomplished by taking one (or multiple exposures) images with the camera each time a LED is turn on. This results in 1080 images acquired per camera position and per exposure.

2.5 First Reconstruction Results

Figure 2.7 presents virtual reconstructions for 8 of the 10 acquired cloths of the MEB collection. These reconstructions are obtained using averaged hemispherical reflectivity (i.e., a color per

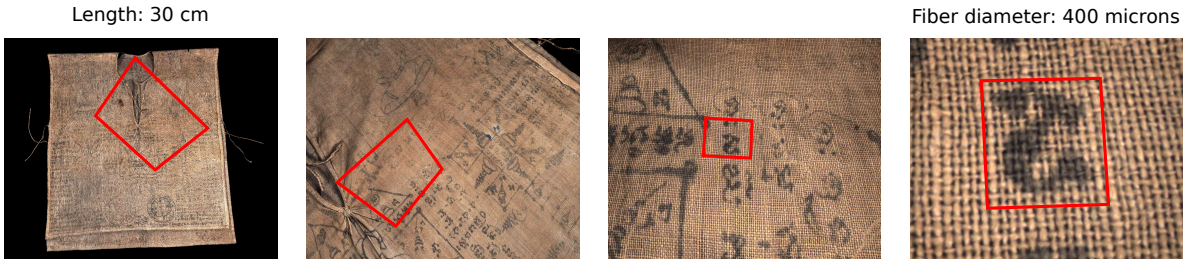


Figure 2.8: Zoom-in images obtained for the reconstruction of the prophylactic shirt.

surface point) thanks to the 4 White LED panels. The number of images per cloth is approximately 20000. In a similar manner to [CDPS09], to assign a color per vertex, we project each image on the mesh (acquires thanks to the 3D scanner) and average (outliers of the color distribution are removed) per vertex a color. However, we discard pixels that are out-of-focus or in the shadows, and we tesselate and refine the mesh in order to align the geometric resolution from the scanner (200 microns) to the optical one (50 microns). As shown in Figure 2.8 these first results already allow some virtual inspections up to 400 microns with calibrated uncertainties (see next Section).

2.6 Calibration

2.6.1 Overview

Applied to *La Coupole*, where we measure a SV-BRDF on a non-planar, two types of calibration must be done:

- **Radiometric** calibration that establishes the amount (i.e., flux) of light reaching the surface and the response curve of the sensor-lens system.
- **Geometric** calibration that establishes the local light and view directions for each point of the object surface.

In the following we only give the main steps of the calibration procedure, the complete details are available in A. Lucat's PhD. manuscript [Luc20].

Radiometric Calibration. The camera radiometric calibration consists in establishing different factors:

(i) the spatial sensitivity of the sensor, which is calibrated by taking multiple photographs of the flat white LED panel. This helps compensating the vignetting effect.

(ii) the curve sensitivity that aims at establishing the relationship between a pixel value (for a given exposure time) and the flux (or another radiometric quantity) received on the pixel.

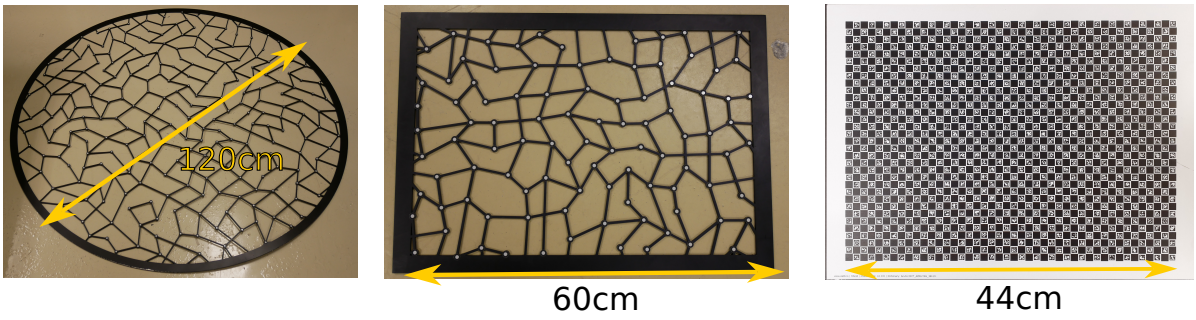


Figure 2.9: **Left and Center:** Calibration targets used with te HandyScan scanner to acquire the geometry of the object. **Right:** Charuco (44 × 33cm) target use to calibrate the positions of the camera, light sources and the object.

(iii) the chromatic correction, which is accomplished using a MacBeth color checker. This allows to establish the correction from the sensor RGB color space to sRGB for example.

Instead of a direct measurement of the LED, our strategy will be [6] to measure the irradiance on each point of the Dibond support. This will be done by a spectralon band of 10 cm width and 70 cm length, which is photographed (one image per LED) and then rotated to cover the whole Dibond support.

Finally, the reader should keep in mind that the radiometric calibrations are done relatively (up to scale factor) and could be done absolutely in the future.

Geometric Calibration. The most important problems to be solved for this calibration step are:

(i) To establish the position of the camera with respect to the robot basis frame.

(ii) To establish the position of each LED with respect to the robot basis frame.

(iii) To establish the rigid transformation (called mesh-to-robot below) between the frame of the mesh, which is obtained from the 3D scanning process, and the robot basis frame.

Once the camera, the LEDs, and the mesh are expressed in the same frame, one can for each pixel of an acquired image, trace a ray to find the intersection with the mesh, and then deduce the view and light directions from the local surface normal.

To achieve steps (i) and (iii), we take 30 pictures of a Charuco calibration target (cf. Figure 2.9). Its fabrication uncertainties are announced to be $100\mu m$. The Charuco target is also used to calibrate for extrinsic and intrinsic camera parameters (effective focal lens, distorsions, focus distance,...) as done usually (cf. Zhang [Zha00]). Contrary to previous approaches, we calibrate intrinsic parameters for each different focusing distances. If a non-calibrated focusing distance is used during the measurement step, we linearly interpolate the intrinsic parameters.

[6] This remains an immediate future work

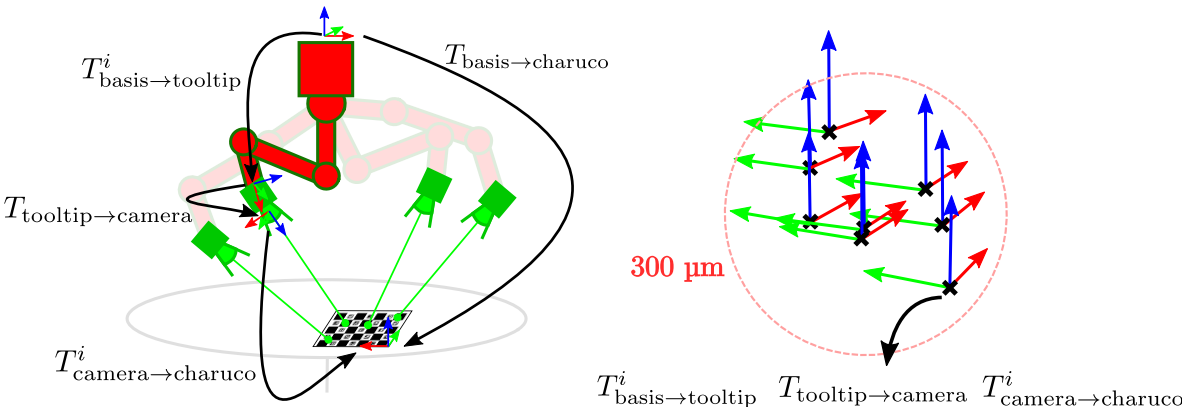


Figure 2.10: Our goal is to infer $T_{\text{tooltip}\to\text{camera}}$, the transformation from the tooltip to the optical center of the camera. **Left:** For each position i, the camera takes an image of the Charuco target. The transformation $T^i_{\text{basis}\to\text{tooltip}}$ (from the robot basis to the robot tooltip) varies for each acquired image. The same applies to $T^i_{\text{camera}\to\text{charuco}}$. However, $T_{\text{basis}\to\text{tooltip}}$ as well as $T_{\text{basis}\to\text{charuco}}$, are identical for all images. **Right:** The final maximum uncertainty of the calibration target position wrt. to the robot basis (i.e., $T_{\text{basis}\to\text{charuco}}$) is 300 microns.

To calibrate the **position of the camera** one must choose a camera model. We use a thin-lens camera model and the calibration procedure goal is to find the position of the optical center.

As shown in Figure 2.10, finding the optical center of the camera requires to know two transformations: one from the basis to the tooltip of the robot and the other from the tooltip to the optical center. When taking a photograph[7] of the Charuco target, we observe that (i) the target never moves wrt. to the robot basis (i.e., $T_{\text{basis}\to\text{charuco}}$ is fixed over time), and (ii) the transformation from the tooltip to the optical center does not change since the camera is firmly [8] attached to the robot. This makes it possible to use the numerical optimization from [Len89] to deduce $T_{\text{tooltip}\to\text{camera}}$ as well as $T_{\text{basis}\to\text{charuco}}$ since:

$$T_{\text{basis}\to\text{charuco}} = T^i_{\text{basis}\to\text{tooltip}} T_{\text{tooltip}\to\text{camera}} T^i_{\text{camera}\to\text{charuco}} \tag{2.1}$$

Note that for each image i, we obtain an estimate of $T_{\text{basis}\to\text{charuco}}$, the final estimate is obtained through averaging all estimates.

The sharper the acquired image the more pixels will be kept and used as measurement samples. The sharpness of an image depends of the depth-of-field which itself depends on the aperture, and thus the circle of confusion of the focus distance. Once the position of the camera is known we calibrate the **focus distance** by taking pictures of the Charuco target from different distances. For each distance, we use the autofocus of the camera to obtain the sharpest image. This procedure establishes a correspondance between requested focus distance and the real one obtained.

[7] At least 3 photographs of the Charuco target are required to insure that the numerical optimization will converge.
[8] If that is not the case (e.g., screws no tighten enough!), the described calibration procedure will detect it!

To establish the **position of each LED** in the robot frame, each LED is photographed (at least 6 times) and a triangulation algorithm is used. This part of the calibration takes three hours but it can skipped (except for the LEDs on the door) as long as *La Coupole* is not moved.

The output of the 3D scanner is a complete mesh where vertices are expressed relatively to an arbitrary frame but in order to facilitate the task, we chose the frame to be on one of the calibration stickers. To establish the *mesh-to-robot* transformation, which will "put" the mesh in the robot frame, we use the algorithm introduced by Sorkine et al. [SR09]. It finds the rigid transformation between the positions of the calibration stickers expressed in the mesh frame, and the ones in the robot frame. The position of the stickers in the robot frame are measured by optical triangulation as done for the LEDs position.

2.6.2 Accuracy and Repeatability Assessment

The assessment of the overall calibration procedure led us to the following metrological results.

Resolution. In real conditions, the resolution of the 3D laser scanner is not $50\mu m$ as announced by the manufacturer but $200\mu m$. Acquisition made with a $50\mu m$ resolution are too noisy to be exploited. Regarding the SV-BRDF, the optical resolution is $50\mu m$ whereas the angular one is 0.7 degree in average. Since the optical resolution is better than the geometrical one, we can average a group of 4×4 pixels to get a coherent evaluation of the SV-BRDF for a surface of at least $200^2 \mu m^2$. Alternatively, one may also tesselate the mesh twice. This allows to have optically spatially varying features assuming that there is no normal variation inside a triangle of $200\mu m$ sidelength.

The average angular resolution comes from:

- estimating the angular resolution of a LED as seen from a point on the object surface. Since the closest (resp. farthest) distance between a point on the surface and a LED is 60cm (resp. 137cm), and since the LED size length is 7mm, the angular resolution is given by $2\arctan\left(\frac{0.007}{2\cdot 0.60}\right) = 0.67\,\text{degree}$ (resp. $0.29\,\text{degree}$).
- estimating the angular resolution (θ_{view}) of the lens aperture as seen from a point on the object surface. For example, assuming a f-number of 8, a 35mm focal length, and a distance of $30\,cm$, we get:

$$\text{lens}_{\text{diameter}} = \frac{0.035}{8}\,\text{m}. \quad \Longrightarrow \theta_{\text{view}} = 2\arctan\left(\frac{\text{lens}_{\text{diameter}}}{2\cdot 0.30}\right) = 0.08\,\text{degree}. \tag{2.2}$$

Note that the worst-case scenario with our Sigma lens is an angular resolution of 0.7 degree, which corresponds to $f\# = 1.4$ and a focusing distance of 20cm. Naturally, we tend to avoid using such a low f-number because the depth of field is too narrow resulting in a big lost of pixels that are out-of-focus.

Uncertainties. The **geometrical** uncertainties are 30 microns (as specified by the manufacturer) for the spatial part (position of a vertex) and 10 degrees for the angular part. The angular part represents the uncertainty of the normal direction deduced from the position of the mesh vertices. This uncertainty is bigger than what we are looking for (0.5 degree);this issue is discussed in Section 2.7.1.

The **optical** uncertainties are 300 microns for the spatial part and 0.1 degree for the angular one. In other words, the camera position is known at 300 microns. Reducing this uncertainty is a future venue as discussed in Section 2.8. Since the calibration of the camera is required to compute the view direction of the BRDF as well as the positions of each LED, reducing the camera position uncertainty is crucial.

2.7 Discussion

2.7.1 Normal Direction Uncertainty

The 3D scanner is calibrated using the vendor procedure. The result is accurate (200 microns resolution for 30 microns of accuracy) since the distances measured by the scanner are linked to the meter standard from the *Institut National de Metrologie*.

However, the normal accuracy is disappointing. This comes from the fact that the smaller the triangle, the higher the accuracy of the vertex positions must be. Indeed the normal direction is computed from the vertex positions. To achieve our targeted accuracy (0.5 degree) one would require a vertex position accuracy of 2.5 microns for a triangle of 200 microns on each side.

On the contrary, if we limit the spatial resolution to 2200 microns, we would obtain the required 0.5 degree accuracy for the normal. This would have for consequence to adjust our appearance representation in order to reproduce details below the 2200 microns resolution. For example, this could be done by encoding per triangle a «micro» spatially-varying BRDF that encodes all light-matter interactions of finer details. Appropriate representations of the geometry and the SV-BRDF to achieve real-time rendering is a future work (cf. Section 2.8) that will be the focus of a PhD thesis or post-doc, starting in 2023, thanks to a Regional grant.

2.7.2 Acquisition Time and Storage Constraints

The acquisition time is related to the size of the object and the desired spatial and angular resolutions, which are themselves limited by the hard drives storage capacity.

The acquisition time depends also on the acquisition speed that the camera can achieve. This is especially true for short exposition times. The theoretical supported speed of the camera is 100fps, which is equivalent of to acquire an image every 10ms. In practice, images acquired with an exposure of less than 50ms are unusable because they overlap each other.

Theoretically speaking the camera framerate (in Hertz) is limited by:

$$1/(t_{\text{exp}} + t_{\text{readout}}), \tag{2.3}$$

where t_{exp} represents the exposure time and t_{readout} the readout time. The readout time, is the time to transfer the content from the camera framebuffer memory to the PC memory. In other words, the readout time is not guaranteed due to the internal electronics of the sensor that introduce additional delays.

To summarize, *La Coupole* generates 20 GB of data per hour when using the hemispherical lighting and 230 GB per hour using the 1080 LEDs. This corresponds to acquire between 100 and 750 view directions per hour using hemispherical lighting, and between 6 and 20 view directions when using LEDs. To be complete, the size of the mesh should also be added to the amount of data generated but it is almost negligible compared to the number of acquired images. Indeed, in our experiments we ended with meshes whose sizes vary from 5GB up to 12 GB whereas as the total size for the images is between 10 and 20 To (for 5 days of acquisition).

2.8 Future Work

We start this section by presenting the most important future work related to the hardware setup, and we finish with some challenges that need to be solved regarding the compression and representation of the acquired data.

2.8.1 Improving *La Coupole*

Geometric accuracy. We note that the 200 microns limitations from the 3D scanner is a limitation that could be lifted with more recent hardware[9] or by combining photogrammetry techniques with scanning technique. In order to combine photogrammetry and laser scanning we would modify the acquisition procedure by intertwining the hemispherical reflectivity and SV-BRDF measurement steps. In other words, after positioning the robot, one photograph is made with full lighting (using the 4 white LED panels) and then for the same position 1080 images are acquired (one for each LED).

Reducing the uncertainties related to the light and view directions is also an interesting venue. This can be done by improving the robot calibration procedure. More precisely, one should not rely on the values provided by the robot to assess the position of the tooltip, but rather (i) use an external tool to determine the potential offset (as done in [Hal07]) of the motors of the robot, or (ii) modify the current optimization procedure (cf. Section 2.6) by introducing additional iteration steps. This would be a 2-step crossed optimization procedure. During the first step the robot motor values are assumed to be correct and a tooltip to camera transformation is optimized. Then, a second optimization step is performed where the tooltip to camera transformation is fixed

[9]The 3D Scanner was bought from Creaform in 2016, the last release promises a 25 microns resolution...

and the motors values are optimized. The procedure is repeated until a user-defined convergence threshold is reached.

Radiometric improvements. A first improvement related to radiometry is to track more precisely the cables that provide power and data transfert for the camera. Indeed, they can block the light and cast shadows on the object, making some pixels unreliable for SV-BRDF measurements. At the moment only the shadows cast by the robot are directly detected by the system.

We try to minimize the effects of diffraction in the images by using an appropriate aperture but we could use our method [3] to discard more precisely pixels affected by the diffraction. This would definitively improve the quality of the measurements at the cost of increasing the post-processing time.

Although, not strictly required, the absolute calibration of the light sources and the camera would be beneficial for the setup. As stated by Palmer and Grand [PG10], metrologists define calibration as the process of assigning absolute engineering units (i.e., radiance, temperature,...) to acquired data (volts, digital counts). In other words, the calibration goal is to determine the transfer function of the instrument. For radiometric quantities, calibration relies on the fact that one can measure electrical current with high accuracy. The requirement to reach an absolute calibration of *La Coupole* are the followings:

- An absolute calibration of the camera sensor using a light source of known power (or one that can be calibrated with an external instrument).
- An absolute calibration of each LED by imaging with the calibrated camera the spectralon sample (absolute BRDF measurement).
- An absolute calibration of the 4 white LED panels by imaging them with the calibrated camera (absolute hemispherical reflectivity measurement).

Speed Acquisition. At the time, when acquiring the appearance of an object without prior, the speed bottleneck is mainly due to the time spent computing a collision-free trajectory for the robot. This non-trivial engineering issue will be solved by using state of the art motion planning library (e.g., MoveIT [CSCC14]). The second identified bottleneck is the computation of the shadow cast by the robot, which could be optimized by using faster visibility query on the GPU (e.g.. Optix [PBD+10]).

Another venue, is to introduce, thanks to external studies from experts, some prior knowledge (e.g., type of fabrics or material used to manufacture it) about the object to pilot dynamically the acquisition process. This would have for consequence to acquire less data and thus reducing the acquisition time.

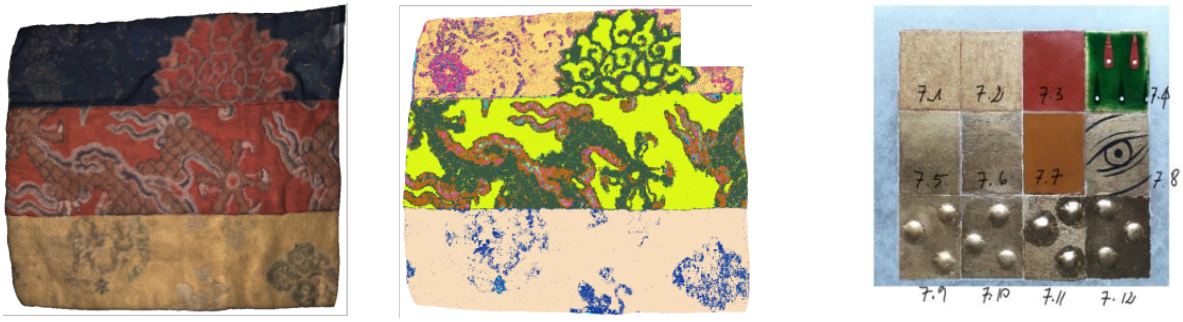


Figure 2.11: **Left:** Color reconstruction obtained thanks to our measurements with the hemispherical lighting. **Center:** Automatic segmentation, based on chromaticity coordinates, of potential BRDFs. **Right:** Material samples manufactured and painted by an artisan on alabaster. Observe that most samples exhibit a gold color but they will probably not exhibit the same BRDF since their underlying structure is different.

Increasing the depth-of-field would improve the acquisition speed since less robot movements would be required. To go behind the limitations of classical lenses, light field cameras, which use an array of micro-lenses for example, could be an interesting upgrade if their spatial resolution (raytrix approach requires a sensor 4 time bigger) and acquisition speed can improve (without exploding prices).

2.8.2 Efficient Representations

Since, terabytes of acquired data cannot fit in the CPU RAM nor GPU VRAM, an important venue for future work is to find suitable representation without loss of quality to visualize the measurements produced by *La Coupole*. This visual lossless compression goal should be guided by the type of applications but also by the hardware used to visualize the representation.

Therefore, a short-term goal is to develop new models, multi-resolution representations as well as visualization solutions in order to render in real time (i.e., with dynamic view and light directions) with a very high fidelity the objects in their virtual form. This type of visualization is intended for experts who want to study the virtual object as well as the general public for real-world or virtual exhibitions

More precisely, our approach is based on the following guidelines:

- An **automatic segmentation method** in the space of the BRDF. As illustrates in Figure 2.11, segmenting only on an average color will not be robust enough. However, it will serve as a first criterion in a hierarchical approach. A second criterion, applied on top of the first one, is to compare view dependent hemispherical reflectivities and then as a last criterion, to compare BRDFs directly. In order to compare directional quantities, we will project them into the Spherical Harmonic basis. Indeed, this will allow us to compare the spherical harmonics coefficients to discriminate between BRDFs. This first step will

permit good compression rate and will serve as a basis for the lowest resolution of our multi-resolution approach.

- A **multi-resolution representation** of the appearance combining both a SV-BRDF, adapted to the observation distance, and a geometric level-of-detail consistent with the SV-BRDF. The objective here is to obtain maximum continuity in the shape and appearance when moving from one resolution to another. This step will permit to obtain virtual renderings in real time without sacrificing the fidelity of the acquisition.

- A **signal approach for SV-BRDF representation**. Unlike BRDF models which must rely on the laws of physics to explain the final appearance, we can also adopt a "signal" or "data" approach in order to compress, interpolate or even extrapolate the data. We aim at compressing a multidimensional signal (bi-spectral, angular and spatial) with, for example, tensor decomposition (e.g., PCA, Block PCA, Clustered PCA...).

Finally, there is also a potential venue to design a measurement strategy based on the final SV-BRDF representation, which could depend and benefit from prior knowledge of the material that needs to be measured.

Part II

Reflectance Models

Since the introduction of the BRDF radiometric definition by Nicodemus [NRH+77], BRDF modeling has been an active research area both in Optics and in Computer Graphics since the seventies with the early development of empirical «shading» models such as the ones proposed by Phong [Pho75] or Blinn [Bli77].

With the generalization of CCD and CMOS sensors, researchers from Computer Graphics and Computer Vision fields have developed image-based devices to measure BRDFs. Compared to classical gonio-reflectometer (e.g., [LBGO15]) coming from the Optics community the these new devices provide much more geometrical measurement configurations. The most emblatic measurement database made with a CCD is the MERL-MIT database introduced by Matusik et al. [MPBM03]. The large amount of materials (100) and light, view configurations ($90 \times 90 \times 180$), made available to everyone, has stimulated researchers:

- To evaluate (e.g., [NDM05]) the BRDF available models,
- To develop BRDF model (e.g., [HTSG91]) that are more physically-based,
- To develop BRDF representations that are more suitable to approximate (e.g., Hemispherical Harmonics [GKPB04]) the measurements.

In this part we introduce two contributions related to the approximations of reflectance measurements and BRDF modeling. The first contribution (cf. Chapter 3) introduces a non-parametric BRDF model based on rational functions, whereas the second one (cf. Chapter 4) introduces a two-scale model, which parameters are all related to physical quantities, and that combines micro-facet and diffractive theories

CHAPTER 3

RATIONAL APPROACH FOR BRDF MEASUREMENTS APPROXIMATION

In this chapter we remind the reader how Rational BRDF [9] are defined and how they are used to approximate BRDF measurements. We conclude this chapter with some recent informal comparisons with other non-parametric representations and outline some potential future work.

3.1 Rational BRDF Definition

A Rational Function (RF) of a finite dimensional vector $\boldsymbol{x}$ of real variables x_i is defined as:

$$r_{n,m}(\boldsymbol{x}) = \frac{p_{n,m}(\boldsymbol{x})}{q_{n,m}(\boldsymbol{x})} = \frac{\sum_{j=0}^{n} p_j b_j(\boldsymbol{x})}{\sum_{k=0}^{m} q_k b_k(\boldsymbol{x})} \tag{3.1}$$

where the $n+1$ (resp. $m+1$) coefficients of the numerator (resp. denominator) are represented by the real numbers p_j (resp. q_k), and where $b_j(\boldsymbol{x})$ and $b_k(\boldsymbol{x})$ are multivariate basis functions. Both $p_{n,m}(\boldsymbol{x})/q_{n,m}(\boldsymbol{x})$ and $\alpha p_{n,m}(\boldsymbol{x})/\alpha q_{n,m}(\boldsymbol{x})$ take the same function values for finite nonzero α, and the coefficients p_j and q_k need only be determined up to a multiplicative constant that can be used to normalize the representation of $r_{n,m}(\boldsymbol{x})$. Therefore $r_{n,m}(\boldsymbol{x})$ has no more than $n+m+1$ **free coefficients**.

As cleverly observed by Romeiro et al. [RVZ08], the materials from the MERL database exhibit almost no variation along the ϕ_d dimension of the Rusinkiewicz's parametrization [Rus98].

Therefore, a Rational BRDF is a 2D rational function that approximates BRDF Measurements:

$$\rho(\boldsymbol{l},\boldsymbol{v}) \approx \rho(\theta_h,\theta_d) \approx r_{n,m}(\theta_h,\theta_d). \tag{3.2}$$

where:

- $\rho(\boldsymbol{l},\boldsymbol{v})$ represents a 4D BRDF measurement for a given light $\boldsymbol{l}$ and view $\boldsymbol{v}$ directions,
- $\boldsymbol{h} = (\theta_h,\phi_h)$ is the half-vector used in the micro-facet theory,
- $\boldsymbol{d} = (\theta_d,\phi_d)$ is the difference vector, which coordinates represent the light direction $\boldsymbol{l}$, in a rotated frame where $\boldsymbol{h}$ defines the north pole direction.

3.2 Fitting Measurements Data with Rational BRDF

To find the coefficients p_j and q_k of Equation 3.1, we rely on a fitting technique developed by Salazar et al. [SCV07].

Assume $s+1$ measured values f_i, each of them located at a vector $\boldsymbol{x}_i$. Let $F_i = [\underline{f_i}, \overline{f_i}]$ represent a real-valued interval placed around each f_i. We would like to find a Rational Function $r_{n,m}(\boldsymbol{x})$ that interpolates these intervals:

$$\forall i = 0,\dots,s \qquad \underline{f_i} \leqslant r_{n,m}(\boldsymbol{x}_i) \leqslant \overline{f_i}, \tag{3.3}$$

and this, for the smallest possible value for $n+m$, with $n+m \leqslant s$ (usually $n+m \ll s$).

It can be shown [SCV07] that a robust solution for $r_{n,m}(\boldsymbol{x}_i)$ can be obtained by solving the following quadratic programming problem $\mathscr{P}(n,m)$:

$$\arg \min_{\boldsymbol{c}\in\mathbb{R}^{n+m+2}} |\boldsymbol{c}|_2 \tag{3.4}$$

$$\text{subject to} \tag{3.5}$$

$$\mathbf{A}^{(j)}_{n,m}\,\boldsymbol{c} - \delta|\mathbf{A}^{(j)}_{n,m}|_2 \geqslant 0, \qquad j = 1,\dots,2s+2 \tag{3.6}$$

(3.7)

$$\text{with } \boldsymbol{c} = (p_0,\dots,p_n,q_0,\dots,q_m)^T \tag{3.8}$$

$$\text{and } \mathbf{A}^{(j)}_{n,m} \text{ denoting the } j\text{-th row of matrix} \tag{3.9}$$

$$\mathbf{A}_{n,m} = \begin{pmatrix} b_0(\boldsymbol{x}_0) & \dots & b_n(\boldsymbol{x}_0) & -\underline{f_0}\,b_0(\boldsymbol{x}_0) & \dots & -\underline{f_0}\,b_m(\boldsymbol{x}_0) \\ \vdots & & \vdots & \vdots & & \vdots \\ b_0(\boldsymbol{x}_s) & \dots & b_n(\boldsymbol{x}_s) & -\underline{f_s}\,b_0(\boldsymbol{x}_s) & \dots & -\underline{f_s}\,b_m(\boldsymbol{x}_s) \\ -b_0(\boldsymbol{x}_0) & \dots & -b_n(\boldsymbol{x}_0) & \overline{f_0}\,b_0(\boldsymbol{x}_0) & \dots & \overline{f_0}\,b_m(\boldsymbol{x}_0) \\ \vdots & & \vdots & \vdots & & \vdots \\ -b_0(\boldsymbol{x}_s) & \dots & -b_n(\boldsymbol{x}_s) & \overline{f_s}\,b_0(\boldsymbol{x}_s) & \dots & \overline{f_s}\,b_m(\boldsymbol{x}_s) \end{pmatrix} \tag{3.10}$$

and where $|\cdot|_2$ denotes the Euclidean norm and $\delta > 0$ is set to be the inverse of the conditional number of the matrix $\mathbf{A}_{n,m}$. Furthermore, if $\mathscr{P}$ has a solution, it is **unique** and the returned function $r_{n,m}$ is **pole-free** at each measured value, i.e., $\forall i = 0,\ldots,s, \quad q_{n,m}(\boldsymbol{x}_i) > 0$. Whether a solution exists for a specific n and m, depends on the relation between the width of the intervals F_i and the values n, m, and s.

In the context of BRDF fitting, we impose additional constraint to enforce the positivity of the found solution:

$$\forall i = 0,\ldots,s \qquad \underline{f_i} \geqslant 0\,. \tag{3.11}$$

Compared to micro-facet analytical models (e.g., [APS00]), which require non-linear optimization techniques (e.g., Levenberg-Marquardt or Sequential Quadratic programming [NDM05]) and only guarantee a local minimum, the involved optimization presented above, is that of a strictly convex function whose global minimum is guaranteed to be reached, provided that a feasible solution exists.

On Numerical Stability. Furthermore, to improve the numerical stability when solving $\mathscr{P}$ it is important to choose orthogonal basis functions [20]. Concretely, Legendre or Chebyshev polynomials are better suited for fitting whereas the monomial basis are more efficient (thanks to Horner factorization) when evaluating the BRDF at the rendering stage.

Solving the Quadratic Problem Efficiently. The speed for the resolution of $\mathscr{P}$ greatly depends on the number $s+1$ of intervals to be interpolated. Our original MATLAB implementation can be described as follows.

To quickly find the couple (n,m) we propose an adaptive procedure that operates on a selected small subset of the whole data. The key observation here is that if the current tested $r_{n,m}(\boldsymbol{x})$ rational function is not a solution for the selected subset, it cannot be a solution for the whole dataset.

Of the given $s+1$ data intervals, only a small number s_0 of intervals are uniformly selected and $r^{(0)}_{n_0,m_0}(\boldsymbol{x})$ is computed such that it satisfies Equation (3.3) for these s_0 data; these s_0 intervals are called the *training data*. Then it is checked how many of the original $s+1$ interval interpolation conditions are automatically satisfied by $r^{(0)}_{n_0,m_0}(\boldsymbol{x})$ in addition to the s_0 imposed conditions. Usually this is quite a lot more. These $s+1-s_0$ intervals are called the *verification data*. Among the violated interval interpolation conditions, we select $s_1 - s_0$ ($s_1 - s_0 = 1$ in our implementation) additional data points to compute $r^{(1)}_{n_1,m_1}(\boldsymbol{x})$ that satisfies Equation (3.3) for these s_1 data. In other words, we update the set of training data. These $s_1 - s_0$ additional training data are selected where $r^{(0)}_{n_0,m_0}(\boldsymbol{x}_i)$ deviates most from the given intervals F_i. With $r^{(1)}_{n_1,m_1}(\boldsymbol{x})$ we then check the new $s+1-s_1$ verification data again. And so on, until the obtained rational model satisfies all verification data.

Material	$(n+1, m+1)$				Maximum degree		
	R	G	B	Total	R	G	B
`beige-fabric`	(18,12)	(21,12)	(18,12)	93	(5,4)	(5,4)	(5,4)
`blue-metallic`	(20,28)	(22,24)	(16,16)	126	(5,6)	(6,6)	(5,5)
`nickel`	(27,16)	(40,9)	(36,12)	140	(6,5)	(8,3)	(7,4)
`chrome`	(84,17)	(47,14)	(50,26)	238	(12,5)	(9,4)	(10,6)

Figure 3.1: Number of coefficients and maximum degrees found for each material approximated with a rational function. The number of coefficients are provided (center column) for the numerator and the denominator each per (R,G,B), as well as (right column) the maximum degree.

More recently, the same kind of strategies were implemented in the C++ ALTA [21] library, which relies on the Quadprog++ [Gue20] library for the rational fitting part. At the moment, the Quadprog++ library[1] implements three different strategies to deal with the large number of inequality constraints of problem $\mathscr{P}$:

- Worst First: Default strategy: check all inequalities, pick the worst one (if any), solve it.
- Worst Set First: Check all inequalities, and pick about `initial_chunk_size` worst ones, and then satisfy them all using the default strategy.
- Sliding Windows : Check the current chunk of `initial_chunk_size` inequalities, pick and solve the worst one (if any), shift to the next chunk and repeat until all chunks are satisfied.

The main advantage of ALTA's C++ implementation is its speed when dealing with problems of high dimensions ($\geqslant 4$).

Fitting Procedure. The fitting procedure is described in Algorithm 1 . As shown in the procedure, the maximum number of coefficients C_{max} is chosen by the user and the algorithm exits as soon as a solution is found. Algorithm 1 may find multiple solutions, but it always selects the more stable solution by choosing the one associated with the lowest condition number of matrix $\mathbf{A}_{n,m}$. If the algorithm does not return a solution, either the number of coefficients is too low or the confidence intervals F_i are too tight and the user must decide to adjust one or the other.

3.3 Results on some Matusik Materials

The most comprehensive and accurate database of measured isotropic BRDFs, known as the *MERL-MIT BRDF database*, combines over 1 billion individual BRDF measurements generated by Matusik et al. [MPBM03]. The measured BRDF data are available with a $90 \times 90 \times 180$ angular

[1]The QuadProg++ library [dGM07, dGM16] was originally developed by Luca du Gaspeo (2007-2010) and Eric Moyer (2009). Later, Gaël Guennebaud modified the library to use the Eigen library and also improved its efficiency by adding the different strategies.

Algorithm 1 Overview of the fitting procedure.

$pdata$ = PREPROCESSED($data$)
C_{max} = RF maximum number of coefficients
$\{F_i\}$ = intervals for measured value $\boldsymbol{x}_i$ of $pdata$
$b_0,...,b_{C_{max}-1}$ are ordered basis functions

FITDATA($pdata$, C_{max}, $\{F_i\}$)

function FITDATA($data$, C_{max}, $\{F_i\}$)
  FITRATFUNC($data$, C_{max}, $\{F_i\}$)
  while `no solution` **do**
    increase the width of all intervals $\{F_i\}$
    `solution` = FITRATFUNC($data$, C_{max}, $\{F_i\}$)
  end while
  return `solution`
end function

function FITRATFUNC($data$, C_{max}, $\{F_i\}$)
  $S = \emptyset \quad S_{best} = \emptyset \quad ccn = \infty$
  for $n = 0,...,C_{max}-1$ **do**
    for $l = n,...,0$ **do**
      cn = CONDITIONNUMBER($\mathbf{A}(l, n-l)$)
      S = SOLVE($\mathscr{P}(l, n-l)$, $data$, $\{F_i\}$)
      if $S \neq \emptyset$ and $cn < ccn$ **then**
        $ccn = cn$
        $S_{best} = S$
      end if
    end for
  end for
  return S_{best}
end function

sampling in $(\theta_h, \theta_d, \phi_d)$, which represents a storage amount of 33 MB per material (the database was last updated in 2006). We focus here on four representatives of common BRDF families: `beige-fabric` (almost perfect Lambertian reflection), `blue-metallic-paint` (glossy reflection with strong chromatic behavior), `nickel` (specular reflection), and `chrome` (almost perfect mirror reflection).

3.3.1 Error Analysis in BRDF space

For each of these materials, we fit a bivariate rational function $r_{n,m}(\theta_h, \theta_d)$ that satisfies Equation (3.3). Of course, the behavior of the approximation is determined by the choice of the intervals $F_i = \left[\underline{f_i}, \overline{f_i}\right]$ containing $\rho(\theta_h^i, \theta_d^i)$, $(i = 0,\ldots,90\times 90-1)$, for each of the BRDF values. Ideally, the widths of these intervals are calibrated such that they respect the accuracy of each individual measurement. Since measurement error bounds are not explicitly available here, the interval widths are taken such that the renderings are visually satisfactory, while leaving a reasonable number $m+n$ of coefficients (cf. Figure 3.1) in the approximation. As a general guideline, we

control the interval widths with the following formula:

$$F_i = [\underline{f_i}, \overline{f_i}] = [\rho(\theta_h^i, \theta_d^i)(1-\epsilon_i), \rho(\theta_h^i, \theta_d^i)(1+\epsilon_i)]$$

where ϵ_i controls the desired relative error.

As shown in Figure 3.3, for all materials we noted increased acquisition noise when approaching grazing angles ($\theta_h > 60^\circ$ and/or $\theta_d > 60^\circ$). In the case of highly specular materials, we also noted that the acquired BRDF values around the center of the lobe ($\theta_h = 0$) are particularly large and subject to noise on `chrome` and `nickel` BRDFs. As shown in Figure 3.2, these acquisition artifacts are also visible when used for rendering. As explained in the first part of this document, lens flare artifacts come from the blades of the diaphragm that generate diffraction. Finally, the noise at the center of the specular lobe comes probably from the saturation of the CCD sensor.

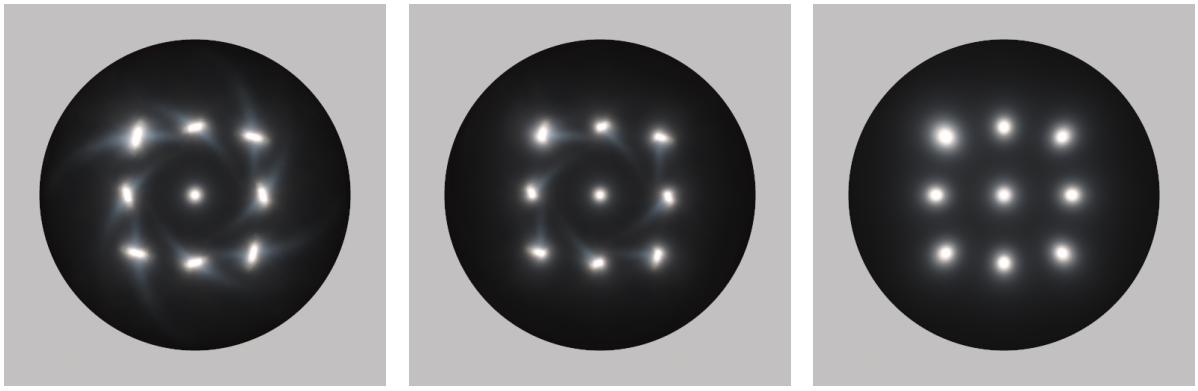


Figure 3.2: Lens-flare-like acquisition artefacts of the `grease- covered-steel` MERL-MIT material illuminated by 9 directional sources. **Left:** Original 3D data. **Middle:** 3D data without noisy grazing angles (i.e., θ_h or $\theta_d > 80^\circ$). **Right:** 2D projection on (θ_h, θ_d) has the effect of filtering most artefacts.

Small interval widths are thus chosen on data with low noise, leading to a very accurate fit in such a configuration. Finally, Figure 3.4 presents the maximum relative errors obtained for BRDFs fitted with rational functions, the micro-facet BRDF model from Ashikhmin et al. [APS00] and polynomial functions (with the same number of coefficients used for rational functions).

3.3.2 Error Analysis in Image space

Figure 3.5 (bottom row) presents the rendering obtained with the Rational BRDF approximation of our four selected materials. The environment map is approximated by 1024 directional light sources selected according to their power in the surrounding environment map. As can be observed, the visual error (measured in Lab color space) is approximately of the same magnitude as with the 2D projected data, but our Rational BRDFs require between 165 and 1000 times less storage memory (in our current implementation, each coefficient p_j and q_k of the RF is stored as an 8-byte double-precision float). When compared to the original 3D data, the compression rate

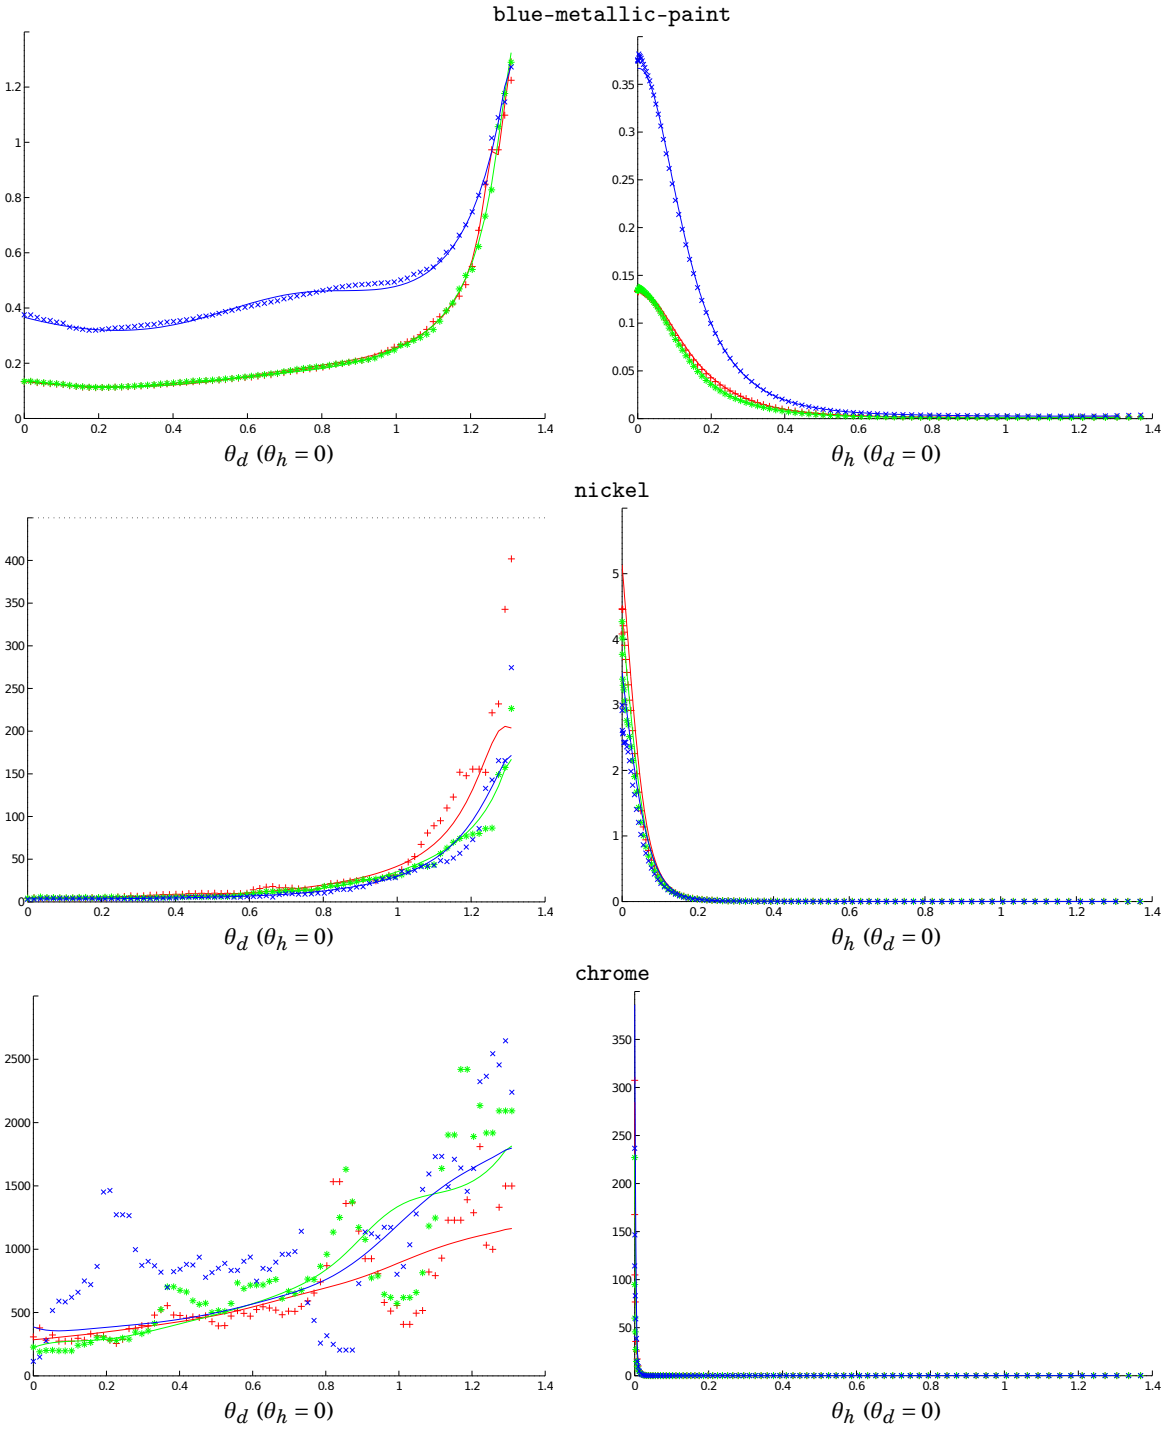


Figure 3.3: Specular materials' BRDF data (dots) without grazing angles (i.e., θ_h or $\theta_d > 75°$) and corresponding rational function approximations (full line) per color channel (RGB).

ranges from 28700:1 to 174000:1. Note that using 4-byte standard floats is likely to be sufficient in many cases. Finally, for grazing angles configuration and for rendering purposes only, we extend

	2D proj. Romeiro 190 KB Error on 3D data	Rational Functions 0.19 KB – 1.15 KB Error on 2D data	 Error on 3D data	Non-linear A&S 0.0625 KB Error on 3D data	Polynomial Functions 0.19 KB – 1.15 KB Error on 3D data
beige-fabric	0.71, 0.73, 0.69	0.71, 0.60, 0.55	0.71, 0.71, 0.69	1.66, 1.75, 1.68	0.71, 0.72, 0.71
blue-metallic	0.80, 0.80, 1.0	0.25, 0.42, 0.29	0.74, 0.71, 1.01	28.49, 34.80, 17.50	43.47, 77.77, 69.02
nickel	1.12, 2.44, 2.95	0.77, 0.78, 0.86	0.77, 2.45, 1.95	4.99, 4.11, 4.64	$4.7 \cdot 10^4$, $4.9 \cdot 10^4$, $1.5 \cdot 10^5$
chrome	5.67, 4.01, 3.99	1.72, 1.97, 3.36	3.29, 2.38, 3.38	65, 134, 148	$2.6 \cdot 10^4$, $1.3 \cdot 10^6$, $1.1 \cdot 10^6$

Figure 3.4: Maximum relative error of the approximation, per (R,G,B) color channel. In the usual order: 2D projected data compared to 3D original data; rational BRDF compared to the 2D projected data, and compared to the 3D original data; non-linear fitting of the Ashikhmin-Shirley analytic model compared to the original data, and least-squares bivariate polynomials compared to the original 3D data. All data at grazing angles ($> 80°$) have been discarded.

our rational functions by clamping light and view directions that exceed 75 degrees. This does not introduce visual artefacts. For all our figures, we used the same clamping for the data and the rational functions in order to illustrate the accuracy of the fitting. Instead of clamping, other extrapolation (linear, cubic) methods can be used to complete the noisy data removed around grazing angles.

3.4 Informal Comparisons with Recent Approaches

Higher compression rates can be achieved with recent parametric models such as the ones from Löw et al. [LKYU12], or Shifted Gamma Distribution [BSH12] or the one [22] presented in the next chapter. However, at the time Rational BRDF was published, the fitting results with the most recent non-linear micro-facet BRDF (Ashikhmin and Shirley model) were below the ones obtained with Rational BRDFs (cf. Table 3.4).

An interesting point to note is that the fitting progress made by parametric models do not only come from the mathematical expressivity of the model but also from the progress of the error metric (e.g., in log-space for Löw et al. [LKYU12]) or even in image space (e.g., [BP20]).

Although, the more recent work of Bagher et al. [BSN16] relies on micro-facet theory, the proposed model is actually a non-parametric approach since its parameters have no physical meaning. For example, the colored shadowing term does not have any physical sense. In their approach, a BRDF (naive model) is stored with 362 coefficients per color channel (vs. 84 for the worst case with rational functions)

Historically, Zernike Polynomials [KDS96] or Spherical Harmonics have been used since the nineties to approximate BRDF measurements. Sillion et al. [SAWG91] or Westin et al. [WAT92] have used classical spherical parametrization but their approaches require between 80 and 133 coefficients for diffuse or glossy materials, and become impractical for high-frequency BRDF.

Since 2011, other non-parametric approaches have been introduced to approximate BRDF.

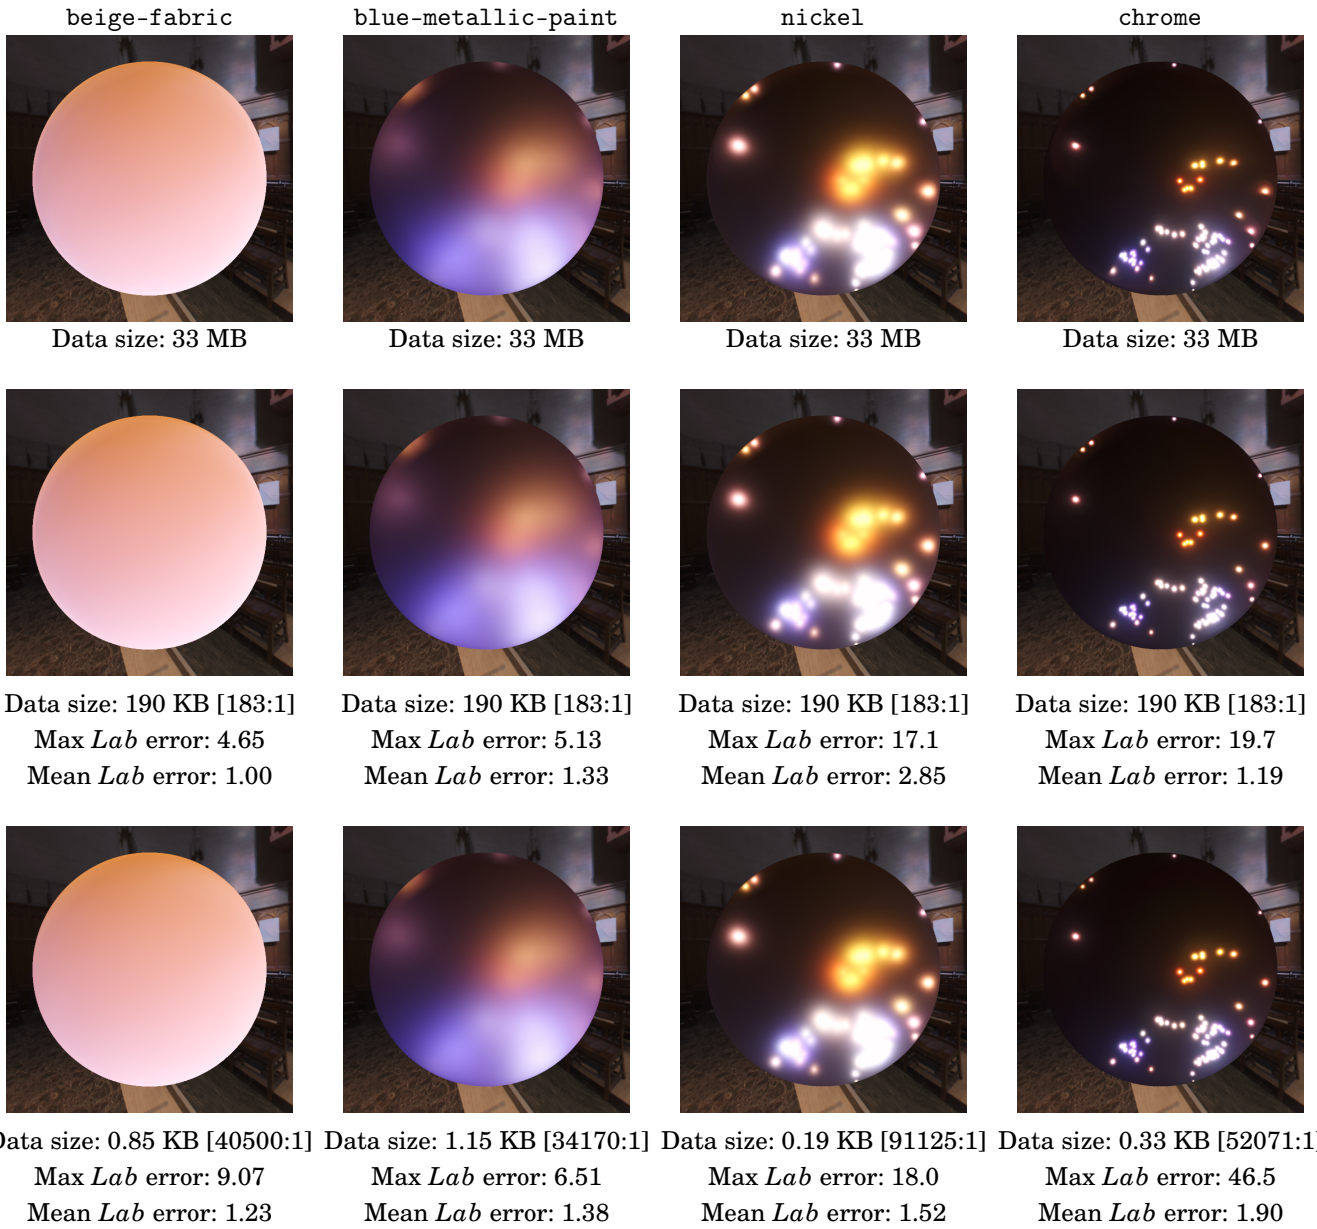


Figure 3.5: From top to bottom: original measured data in $(\theta_h, \theta_d, \phi_d)$, projected data in (θ_h, θ_d), rational BRDF fit in (θ_h, θ_d). All scenes are rendered using 1024 directional light sources sampled from the surrounding environment map. Each Rational BRDF generates a similar magnitude of Lab visual errors as the ones produced by the 2D projected data, but requires between 165 and 1000 times less storage. Note that the data size are provided using float point numbers encoded with double precision (i.e., 8 bytes on classical 64bits CPU).

Recently, the usage of Zonal Harmonics by Soler et al. [SBN15] has been used primarily for Precomptured Radiance Transfer. In fact, they do not achieve higher compression rates than Rational BRDF. Indeed, Cyril Soler reports [Sol21] **per color channel**, 120 coefficients for all

fabric materials, 960 coefficients for most materials of the MERL database and for extreme specular materials, such as the chrome or chrome-steel, $61^2 = 3721$ coefficients.

Finally, tensor approaches, introduced to the Computer Graphics community, by Bigli et al. [BOK11] and improved by Tongbuasirilai et al. [TU+20] are promising but their results show that BRDF are approximated with 226 coefficients in the best case scenario and 360 coefficients in the worst case.

3.5 Future Work

A first avenue to explore as future work is to improve the fitting technique for rational functions, especially for ill-conditioned problems. Indeed, when solving the quadratic programming problem of equation $\mathscr{P}$, the found solution might have a denominator equal to zero on the domain of interest. This comes from the fact that Equation 3.11 is a discreet condition, not a continuous one. In practice, one may always add additional constraints to enforce the positivity of the denominator by interpolating the data. However, as explained, in Section 3.2, additional constraints increase the fitting time. Therefore, a possible solution to investigate could be to use Bernstein polynomials[2], as basis functions in Equation 3.1 with a positivity constraint on the weights. This constraint is linear and easy to add in any quadratic programming solver and will enforce a strictly positive dominator. However, the expressivity of such rational functions might be too narrow to achieve high-compression rates.

A second possible avenue is to combine non-linear fitting technique for rational functions with the vertical segment approach when no confidence intervals are available. Indeed, a rational function is a non-linear function that can be fitted with Levenberg-Marquardt or Sequential Quadratic Programming algorithms. Furthermore, if the global convergence is required, one may use the errors and the solution returned by the non-linear fitting pass, as starting point for the fitting algorithm 1. The errors are then used to setup the interval F_i widths.

Apart from improving the fitting technique, another possible future work is to use 1D or 2D rational functions to compress tabulated data obtained when decomposing a 3D or 4D BRDF. The BRDF decomposition can be obtained with tensorial approaches but also from the homomorphic factorization from McCool [MAA01]. Despite the work from [TU+20] and ours [23], determining the best parametrization to perform the decomposition remains an open question. Decomposition is a research direction also interesting for material editing but is tightly connected to the parametrization. Ideally, for editing purposes, one would like to separate macroscopic shading effects (e.g., specular, diffuse, Fresnel, asperity scattering, retro-reflection, etc.) on each axis of a new parametrization for BRDF.

Finally, as shown at I3D [24], thanks to the Horner factorization, Rational Functions expressed in the monomial basis are very fast to evaluate on CPU as well as on GPU. Therefore,

[2]Bernstein polynomials of degree n are always positive or null and more importantly they form a basis.

they are a viable alternative even for real-time rendering. The main issue at the moment to use them in that context, is the lack of method to pre-convolve with an environment map with a Rational BRDF.

CHAPTER 4

COMBINING MICROFACET AND DIFFRACTION THEORIES

This chapter presents a two-scale BRDF model [22] that combines the microfacet theory [BS87] and a particular diffraction theory called, Generalized Harvey Shack [Kry06]. This model falls into the family of parametric models and moreover, its parameters are all related to physical quantities.

4.1 Introduction

In 2005, Ngan et al.[NDM05] published an extensive study on the capabilities of BRDF models to represent the BRDF measurements of the MERL-MIT database [MPBM03]. They fitted, with the Matlab implementation [Mat22] of Sequential Quadratic Programming (SQP, see Chapter 18 of Nocedal and Wright [JS06] for an introduction), the 100 materials of the dataset with 7 analytical BRDF models. As shown in Figure 4.1, the study concludes that the best BRDF models were the Cook-Torrance [CT82] and a simplified[1] version of the model of He et al. [HTSG91].

The study also shows that adding a second lobe reduces the fitting error and subsequent work [Bur12, LKYU12] did use to a second lobe. The additional lobe can be a Lambertian lobe to take into account subsurface scattering or a generalized cosine term such as the one introduced by Lafortune et al. [LFTG97] or a even a Cook-Torrance one. Obviously, adding a lobe gives more flexibility in terms of mathematical expressiveness but this solution does not make any physical sense, since there cannot be two normal distributions for a single micro-geometry.

By comparing carefully the shape of the measured lobe against the one obtained with the Cook-Torrance model (cf. Figure 4.2), we observe that microfacet-based BRDF models are very good at fitting the large values (lobe center) but remain inaccurate for low values.

[1] Since Ngan et al. did not provide their implementation, it is quite hard to assess precisely what simplifications were made to the model. However, they did point out that they did not use the polarized version of the model.

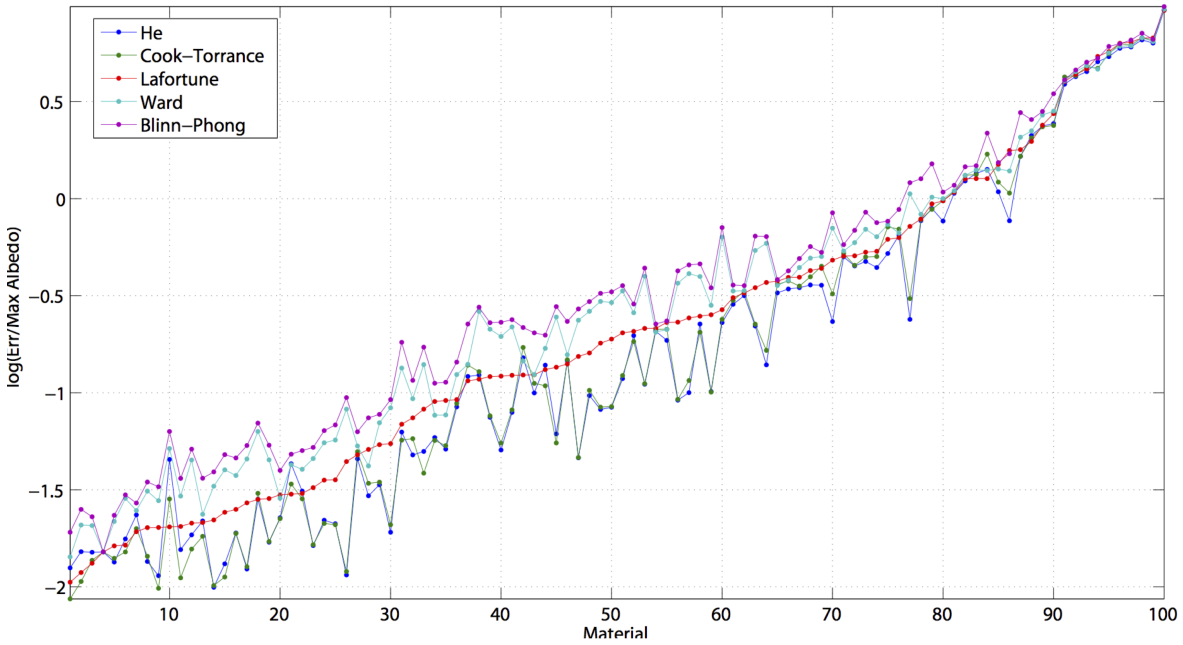


Figure 4.1: Complete fitting results obtained by Ngan et al. [NDM05]. According to the chosen fitting metric, the best results were obtained for the Cook-Torrance [CT82] and the He et al. [HTSG91] BRDF models.

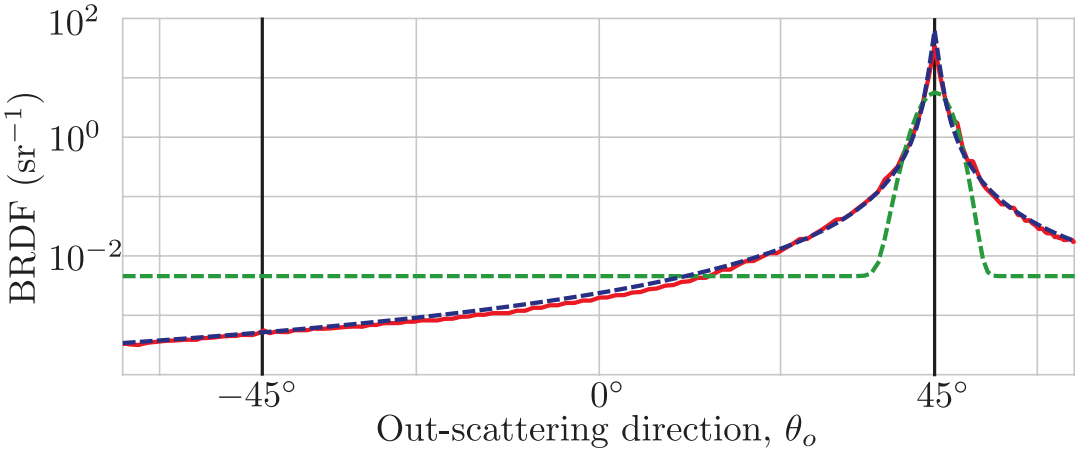


Figure 4.2: .

From Löw et al. [LKYU12]. The red curve represents the blue-metallic paint material measurements for an incident light angle of 45^o. The green curve represents the Cook-Torrance approximation whereas the blue curve represents the *smooth* model introduced by Löw et al. This plot illustrates the fact that a Gaussian function is not adequate to fit low values of the BRDF.

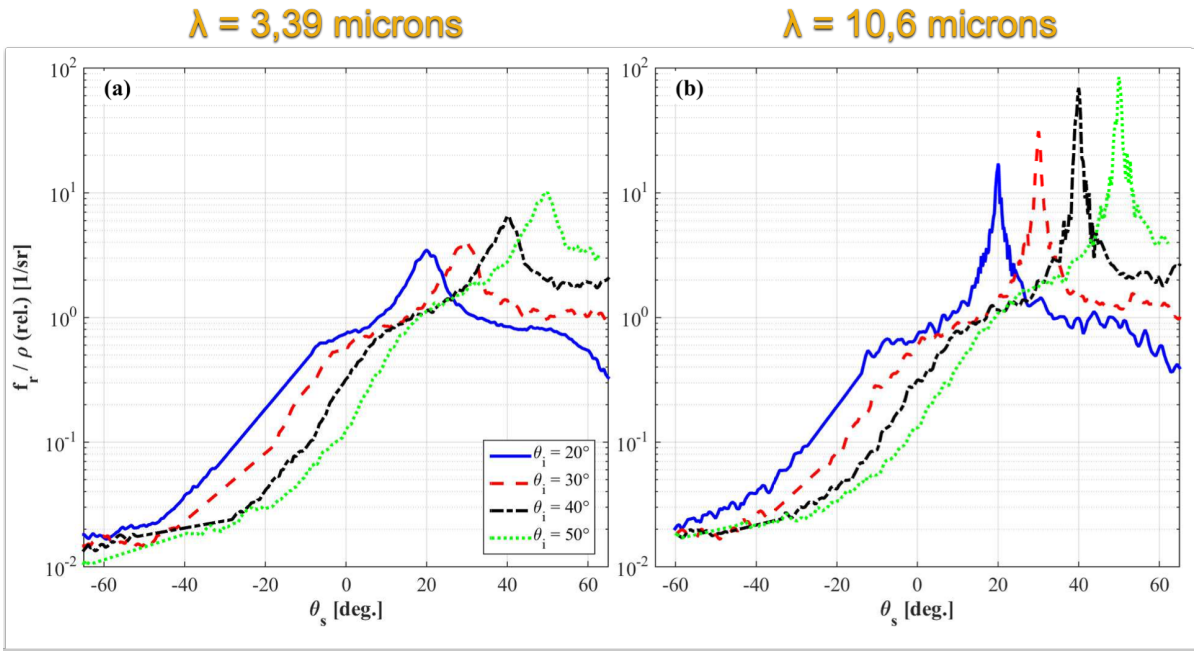


Figure 4.3: BRDF measurements of a grit-blasted nickels from Butler et al. [BNM15]. These measurements illustrate the fact that the shape and the maximum values of the specular lobe vary with the light wavelength. This type of behavior cannot be predicted by the microfacet theory.

A second observation (cf. Figure 4.3) we made in 2016, is that the specular lobe width varies with respect to the light wavelength (λ). This contradicts the idea that a single geometry can capture the entire material response [BNM15].

All these observations lead us to investigate if diffraction phenomena could play a significant role and if by taking them into account, we could develop a more precise and more suitable model for BRDF measurements approximation.

Löw et al. were the first to introduce diffraction models from the Optics community to the Computer Graphics community. Inspired by the Rayleigh-Rice [Ric51a] (see also Stover [Sto95] for an introduction) and the Harvey Shack [Har75] theories, they introduced two BRDF models that are reciprocal but do not necessarily conserve energy. Both BRDF models utilize a power inverse law (the so-called ABC-model) to represent the shape of the specular lobe. With a dedicated fitting technique, they showed that their models perform better than Gaussian shape microfacet functions when it comes to approximate the MERL-MIT database. Although their models are inspired by diffraction theories, Low et al. completely removed the physical parameters of the original diffraction models. The wavelength dependency is absent, the generalized polarization factor, Q, is simplified to the classical Fresnel term.

In view of the above remarks, our main goal was to provide a model that reduces further fitting errors but also give a more plausible physical explanations than the previous approaches.

4.2 Overview

The presented model is the third iteration of two previous models published as INRIA technical reports (cf. [25]).

The two previous models combine empirically two lobes, one from the microfacet theory and another one from the diffraction theory, whereas in this chapter, we make the physical assumption that microfacets are not only Fresnel reflectors but also diffracting elements that produce interferences. From this hypothesis, we derive a complete two-scale model (cf. Section 4.3) and extend it for multiple layers (cf. Section 4.6). For both models we present fitting and rendering results and we conclude this chapter with an informal comparison with the model of He et al. [HTSG91] before outlining some potential future work.

4.3 A Two-Scale BRDF Model

We remind the original microfacet framework as formulated by Walter et al. [WMLT07]:

$$\rho(\boldsymbol{i},\boldsymbol{o}) = \int_\Omega \left|\frac{\boldsymbol{i}\cdot\boldsymbol{m}}{\boldsymbol{i}\cdot\boldsymbol{n}}\right| f_s(\boldsymbol{i},\boldsymbol{o},\boldsymbol{m}) \left|\frac{\boldsymbol{o}\cdot\boldsymbol{m}}{\boldsymbol{o}\cdot\boldsymbol{n}}\right| G_{\boldsymbol{m}}(\boldsymbol{i},\boldsymbol{o}) D(\boldsymbol{m})\, \mathrm{d}\omega_m. \tag{4.1}$$

where the macro-surface has normal $\boldsymbol{n}$; each individual microfacet, with reflectance f_s, has normal $\boldsymbol{m}$ following a normal distribution $D(\boldsymbol{m})$. The shadowing-masking term G expresses the probability for the light to be blocked before or after the specular reflection on the microfacet. It is usually computed using Smith's method through a double integration of D [Smi67, WMLT07, Hei14]. For a fixed light ($\boldsymbol{i}$) and view ($\boldsymbol{o}$) directions, the overall reflectance of the surface is the integral of the contributions from each microfacet, averaged by the probability of each microfacet orientation.

The microfacet framework is generic enough to work with any kind of reflectance as long as f_s is a BRDF that satisfies physical constraints (positivity and energy conservation). In the generic case, there is no closed-form solution to this equation. Most previous approaches have focused on the distribution term D and its associated shadowing-masking term G. The microfacet framework is based on the hypothesis that the surface details are much larger than the visible wavelengths of lights so that the rules of geometrical optics apply.

When a microfacet is assumed to be a mirror (perfect or not), $f_s(i,o,m)$ becomes a Fresnel-weighted Dirac:

$$f_s(\boldsymbol{i},\boldsymbol{o},\boldsymbol{m}) = F(\eta(\lambda),\boldsymbol{i},\boldsymbol{o}) \frac{\delta(\mathrm{refl}(\boldsymbol{i}),\boldsymbol{o})}{\cos\theta_o}, \tag{4.2}$$

and the integral 4.1 collapses into a function evaluation that simplifies to the standard and well-known Cook-Torrance form:

$$\rho_{CT}(\boldsymbol{i},\boldsymbol{o}) = \frac{F(\eta(\lambda),\theta_d) D(\theta_h) G(\boldsymbol{i},\boldsymbol{o})}{4\cos\theta_i \cos\theta_o} \quad \text{with} \quad \begin{cases} \cos\theta_i = \boldsymbol{i}\cdot\boldsymbol{n} & \cos\theta_o = \boldsymbol{o}\cdot\boldsymbol{n} \\ \cos\theta_h = \boldsymbol{m}\cdot\boldsymbol{n} & \cos\theta_d = \boldsymbol{i}\cdot\boldsymbol{m} \end{cases}. \tag{4.3}$$

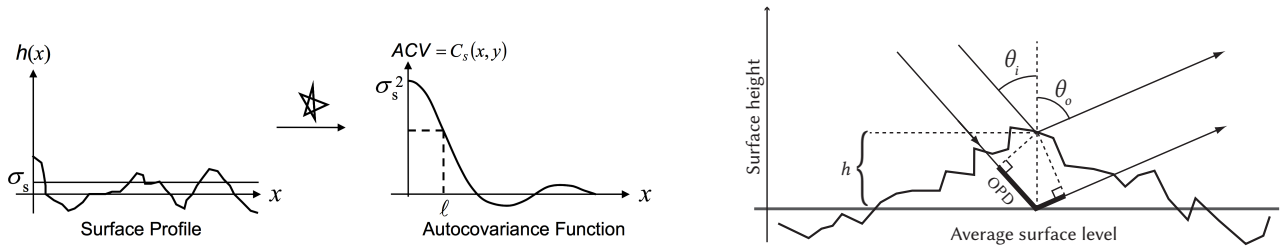


Figure 4.4: **Left, Statistical Parameters of a Surface:**. A 1D surface profile is a distribution $h(x)$ of height along a single axis. If $h(x)$ is Gaussian, then σ_s is the root mean square surface roughness and is the standard deviation of $h(x)$. Furthermore, σ_s^2 is the surface variance. From the 1D surface profile, one may compute the autocovariance function that represents how much the heights of two points are *correlated*. Its maximum value appears when the two points coincide. The autocovariance function at the origin is also equal to σ_s^2. The autocovariance length, l, is defined as the half-width of the autocovariance function at the $1/e$ height. **Right, OPD:**. Light reflected off a surface at a height h has traveled a shorter distance than light reflected on the average surface level. This Optical Path Difference (OPD) causes interference between parallel rays and is responsible for diffraction effects. Here, $OPD = (\cos\theta_i + \cos\theta_o)h(x,y)$ [Kry06]. Heights h are relative to the average surface level.

Looking at equation 4.3 it is obvious that D and G only depend on the *geometry* of microfacets, not on the wavelength of the incoming light. Reflection lobes for the different wavelengths are scaled version of each other, with the *same relative width*. The only wavelength dependency comes from the Fresnel term $F(\eta(\lambda))$, which is responsible for the material color via the index of refraction $\eta(\lambda)$.

4.3.1 A diffractive Microfacet

We now assume that a microfacet is not perfectly smooth and might contain subwavelength imperfections. Its optical response can no longer be modeled by geometrical optics and we need to rely on diffractive optics. A very large literature (cf. [EG04] for a survey) treats the subject of diffraction induced by a non-smooth surface lit by an electromagnetic wave. In this chapter, we consider surfaces that can be described with statistical quantities (e.g., roughness, height correlation, cf. Figure 4.4 left) and we disregard periodic surfaces such as gratings. Due to their periodicity, diffraction gratings also generate interferences but this specific problem can be treated differently (e.g., Stam [Sta99]).

Diffraction BRDF models can be classified depending on the diffraction theory they rely on:

- The Rayleigh-Rice vector perturbation theory ([Ray07, Ric51a]), which is valid only for smooth surfaces and wide-angle scattering (see Stover [Sto95] for a comprehensive overview on the subject).

- The Beckmann-Kirchhoff [BS87] theory and its improvement, named the Modified Beckmann-Kirchhoff (cf. [VH98]) is valid for a wide range of surface roughness and also wide-angle

scattering. However, these theories are restricted to surface with Gaussian profile and autocovariance function.

- The Harvey-Shack [Har75] theory was improved Krywonos [Kry06] who establishes the Generalized Harvey-Shack (GHS) theory. GHS is valid for all type of surface roughnesses and scattering angles, and assumes only that the surface profile is Gaussian. The autocovariance function is *not* restricted to Gaussian profile.

The optical smoothness/roughness mentioned above is always defined relatively to light wavelength considered:

Smooth:	$g \ll 1$
Intermediate:	$g \approx 1$
Rough:	$g \gg 1 \quad \text{with} g = \frac{\sigma_s}{\lambda}$.

For our two-scale model, we decided to rely on the Generalized Harvey-Shack theory as it is applicable to a wide range of configurations and provides a good approximation to measured reflectance. The theory is based on the distance traveled by the light as it is reflected on the surface. Since the surface is not flat, light reflected at a point of height $h(x,y)$ (see Figure 4.4 right) has traveled a shorter distance than if it was reflected at the reference height [Kry06, p. 163]. This difference in optical path length (OPD):

$$OPD = (\cos\theta_i + \cos\theta_o)h(x,y)$$

results in phase variations at the reflected wavefront of $(2\pi/\lambda)(\cos\theta_i + \cos\theta_o)h(x,y)$. Averaging these phase variations gives the reflectance function.

Reflectance $\rho_{\text{diffr.}}$ caused by surface variations is separated in two different lobes: a Dirac in the specular reflection direction, surrounded by a halo of scattered light [Har75, Kry06]. Their relative intensity depends on the surface roughness σ_s^2, defined as the variance of the height distribution h:

$$\rho_{\text{diffr.}}(\boldsymbol{i},\boldsymbol{o}) = AF(\boldsymbol{i},\boldsymbol{o})\frac{\delta(\text{refl}(\boldsymbol{i}),\boldsymbol{o})}{\cos\theta_o} + (1-A)Q(\boldsymbol{i},\boldsymbol{o})S_{HS}(\boldsymbol{f}) \tag{4.4}$$

$$\text{with } A = \mathrm{e}^{-\left(2\pi\frac{\sigma_s}{\lambda}(\cos\theta_i+\cos\theta_o)\right)^2}, \tag{4.5}$$

where S_{HS} is a scattering function (detailed later) , F is the Fresnel term for reflectance and Q is the polarization factor (cf. Appendix 4.A).

As shown in Equation 4.4, when the diffraction lobe increases, the specular peak decreases. Their relative importance depends on surface roughness and directions: at grazing angles, $\cos\theta_i \approx 0$ and A is almost equal to 1; most of the energy goes to the specular reflection. The influence of diffraction is stronger at normal incidence, where $\cos\theta_i \approx 1$.

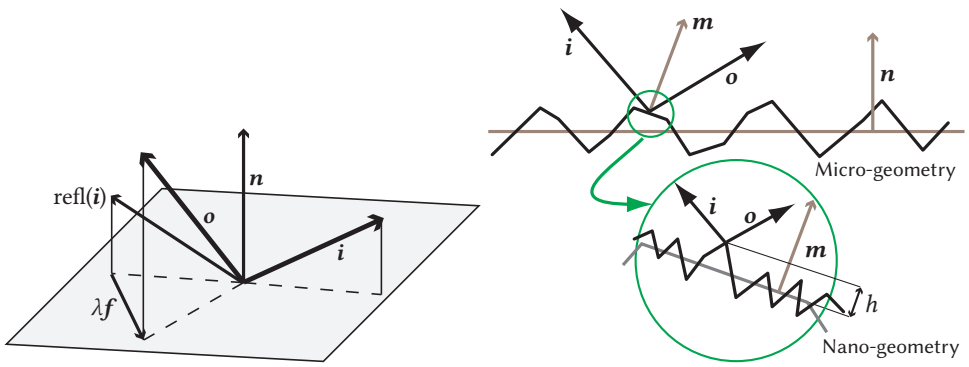


Figure 4.5: **Left:** Diffraction effects depend on $\boldsymbol{f}$, the projected vector between the reflected incoming direction refl($\boldsymbol{i}$) and the outgoing direction $\boldsymbol{o}$, divided by the wavelength λ.**Right:** In our two-scale model, we make the assumption that each micro-facet contains nano roughnesses that are small compared to the light wavelength and therefore diffract the light.

The main parameter controlling the reflectance lobe $\rho_{\text{diffr.}}$ is the 2D vector $\boldsymbol{f}$, built from the projections on the tangent plane of the difference between the specular reflection of the incoming direction and the outgoing direction, divided by the wavelength (see Figure 4.5):

$$\boldsymbol{f} = \frac{1}{\lambda}\text{proj}(\text{refl}(\boldsymbol{i}) - \boldsymbol{o}) = \frac{1}{\lambda}((\boldsymbol{i}+\boldsymbol{o})\times\boldsymbol{n})\times\boldsymbol{n} \tag{4.6}$$

$$f = \|\boldsymbol{f}\| = \frac{2}{\lambda}\sin\theta_h\cos\theta_d. \tag{4.7}$$

By applying $\rho_{\text{diffr.}}$ to each microfacet we can derive the macroscopic BRDF (cf. Figure 4.5 right) by substituting Equation 4.4 into the Equation 4.1, we get:

(4.8)
$$\rho(\boldsymbol{i},\boldsymbol{o},\boldsymbol{n}) = \int_\Omega \left|\frac{\boldsymbol{i}\cdot\boldsymbol{m}}{\boldsymbol{i}\cdot\boldsymbol{n}}\right| \left|\frac{\boldsymbol{o}\cdot\boldsymbol{m}}{\boldsymbol{o}\cdot\boldsymbol{n}}\right| G(\boldsymbol{i},\boldsymbol{o},\boldsymbol{m})D(\boldsymbol{m})\left(A(\boldsymbol{i},\boldsymbol{o},\boldsymbol{m})F(\boldsymbol{i},\boldsymbol{o},\boldsymbol{m})\frac{\delta(\text{refl}(\boldsymbol{i}),\boldsymbol{o})}{|\,\boldsymbol{o}\cdot\boldsymbol{m}|} + (1-A(\boldsymbol{i},\boldsymbol{o},\boldsymbol{m}))Q(\boldsymbol{i},\boldsymbol{o},\boldsymbol{m})S_{HS}(\boldsymbol{f})\right) \mathrm{d}\omega_m.$$

(4.9)
$$= \underbrace{\int_\Omega \left|\frac{\boldsymbol{i}\cdot\boldsymbol{m}}{\boldsymbol{i}\cdot\boldsymbol{n}}\right| \left|\frac{\boldsymbol{o}\cdot\boldsymbol{m}}{\boldsymbol{o}\cdot\boldsymbol{n}}\right| G(\boldsymbol{i},\boldsymbol{o})D(\boldsymbol{m})A(\boldsymbol{i},\boldsymbol{o},\boldsymbol{m})F(\boldsymbol{i},\boldsymbol{o},\boldsymbol{m})\frac{\delta(\text{refl}(\boldsymbol{i}),\boldsymbol{o})}{|\,\boldsymbol{o}\cdot\boldsymbol{m}|}\,\mathrm{d}\omega_m}_{\rho_{\text{microfacet}}}$$

(4.10)
$$+ \underbrace{\int_\Omega \left|\frac{\boldsymbol{i}\cdot\boldsymbol{m}}{\boldsymbol{i}\cdot\boldsymbol{n}}\right| \left|\frac{\boldsymbol{o}\cdot\boldsymbol{m}}{\boldsymbol{o}\cdot\boldsymbol{n}}\right| G(\boldsymbol{i},\boldsymbol{o},\boldsymbol{m})D(\boldsymbol{m})(1-A(\boldsymbol{i},\boldsymbol{o},\boldsymbol{m}))Q(\boldsymbol{i},\boldsymbol{o},\boldsymbol{m})S_{HS}(\boldsymbol{f})\,\mathrm{d}\omega_m}_{\rho_{\text{CTD}}}$$

Note that the quantities from Equation 4.4 are now defined relatively to the microfacet normal m. In another words, the factor A is defined by:

$$A(\boldsymbol{i},\boldsymbol{o},\boldsymbol{m}) = \mathrm{e}^{-\left(2\pi\frac{\sigma_s}{\lambda}(\boldsymbol{i}\cdot\boldsymbol{m}+\boldsymbol{o}\cdot\boldsymbol{m})\right)^2} \tag{4.11}$$

Equations 4.9 and 4.10 show the two lobe components of our two-scale model. Equation 4.10 represents the diffraction lobe (cf. Section 4.3.3) of our two-scale model. Computing efficiently

ρ_{CTD}, the Cook-Torrance Diffraction (CTD) lobe, is our core theoretical contribution. Our technical contribution is to demonstrate better fitting results when combining the CTD lobe and the microfacet lobe (cf. Equation 4.9), which is detailed in the next section.

4.3.2 The Microfacet lobe of the two-scale model

Since equation 4.9 contains a Dirac distribution the only case where the integral is different from zero is when $\text{refl}(\boldsymbol{i}) = \boldsymbol{o}$, or $\boldsymbol{m} = \boldsymbol{h}$. Consequently, the term A becomes equal to:

$$A(\boldsymbol{i},\boldsymbol{o},\boldsymbol{m}) = \mathrm{e}^{-\left(2\pi\frac{\sigma_s}{\lambda}(\boldsymbol{o}\cdot\boldsymbol{m}+\boldsymbol{o}\cdot\boldsymbol{m})\right)^2} = \mathrm{e}^{-\left(2\pi\frac{\sigma_s}{\lambda}(2\cos\theta_d)\right)} \tag{4.12}$$

with $\cos\theta_d = \boldsymbol{o}\cdot\boldsymbol{h}$ and $\boldsymbol{m} = \boldsymbol{h} = (\boldsymbol{i}+\boldsymbol{o})/|\boldsymbol{i}+\boldsymbol{o}|$ at the macroscopic level. Consequently the equation $\rho_{\text{microfacet}}$ 4.9 simplifies to:

$$\rho_{\text{microfacet}}(\boldsymbol{i},\boldsymbol{o}) = \underbrace{\mathrm{e}^{-\left(2\pi\frac{\sigma_s}{\lambda}(2\cos\theta_d)\right)}}_{A_{\text{m}}(\cos\theta_d)}\rho_{\text{CT}}(\boldsymbol{i},\boldsymbol{o}) \tag{4.13}$$

where $\rho_{\text{CT}}(\boldsymbol{i},\boldsymbol{o})$ is defined in equation 4.3.

Instead of using typical distribution terms (e.g., Beckmann or GGX) for $D(\boldsymbol{m})$, we decided to use the Exponential Power Distribution (EPD) introduced to the Computer Graphics community by Brady et al. [BLPW14]. It provides a good fit with the sharp decrease near the specular peak observed on many measured BRDFs. The EPD is controlled by two parameters (cf. Figure 4.6): β, which controls the width of the peak and p, which controls the kurtosis (i.e., whether or not the lobe drops sharply at the origin). The EPD is defined as follows:

$$D_{EPD}(\theta_h) = \frac{\chi_{[0,\frac{\pi}{2}]}(\theta_h)}{\cos^4\theta_h}\frac{p}{\pi\beta^2\Gamma(1/p)}\mathrm{e}^{-\left(\frac{\tan^2\theta_h}{\beta^2}\right)^p} \tag{4.14}$$

To compute the associated shadowing term $G(\boldsymbol{i},\boldsymbol{o},\boldsymbol{m})$, we use Smith's method (i.e., $G(\boldsymbol{i},\boldsymbol{o},\boldsymbol{m}) = G_1(\boldsymbol{i},\boldsymbol{m})G_1(\boldsymbol{o},\boldsymbol{m})$) where

$$G_1(\theta) = \frac{1}{1+\Lambda(\beta\tan\theta)} \quad \text{with:} \tag{4.15}$$

$$\Lambda(\beta\tan\theta) = \int_{1/\beta\tan\theta}^{\infty}(r\beta\tan\theta-1)P_2(r)\,\mathrm{d}r \qquad \text{and} \quad P_2(r) = \frac{p}{\pi\Gamma(1/p)}\int_0^{\infty}\mathrm{e}^{-(r^2+q^2)^p}\,\mathrm{d}q\,. \tag{4.16}$$

where:

1. $\cos(\theta) = \boldsymbol{i}\cdot\boldsymbol{m}$ when computing $G_1(\boldsymbol{i},\boldsymbol{m})$ and $\cos\theta = \boldsymbol{o}\cdot\boldsymbol{m}$ when computing $G_1(\boldsymbol{o},\boldsymbol{m})$
2. $\Gamma(x)$ is the Gamma function.

Since the distribution is shape-invariant, Λ can be expressed as a function of $\beta\tan\theta$ and p. We precompute Λ and store its values in a 2D array indexed by $\beta\tan\theta$ and p (a single 2D array for all values of the three parameters β, p, and θ requires 390 KB of memory). This 2D table could probably be further compressed using rational functions.

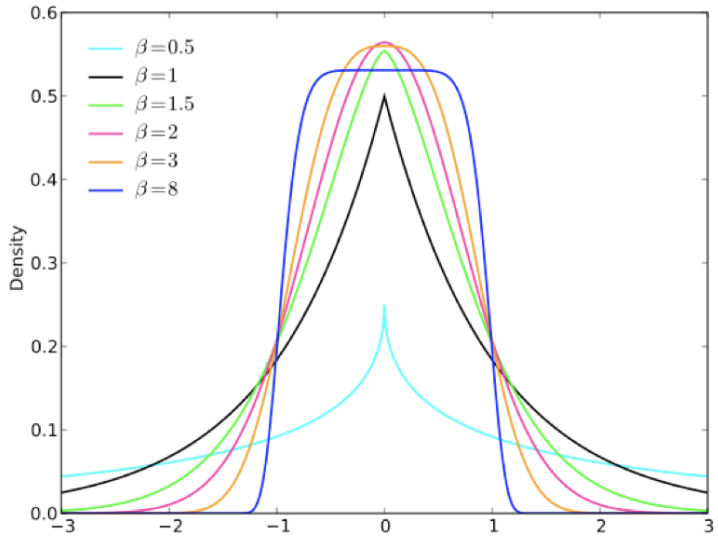


Figure 4.6: The exponential power distribution (EPD) is controlled by two parameters β and p.

4.3.3 The Diffraction lobe of the two-scale model

Before explaining how to tackle the integral from Equation 4.10, two preliminary remarks must be made:

1. Equation 4.10 is solved for the EPD microfacet distribution.

2. The scattering function, $S_{\text{HS}}(f)$ needs to be defined.

Scattering Function $S_{\text{HS}}(f)$. Harvey and Krywonos derived different expressions for $S_{\text{HS}}(f)$ for each type of optical smoothness. However, in our approach, we assume that the diffraction induces by microfacet are due to very small surface variations (i.e., $\sigma_s < \lambda$). This allow us to use an analytical and simplified form for the scattering function (see below) instead of the the generic form that contains a costly Fourier transform for each view and light directions. Note that, large (i.e., $\sigma_s > \lambda$) roughnesses, are modeled by the microfacet lobe $\rho_{\text{microfacet}}$.

In general, the scattering function S_{HS} is expressed with a scaled version of the Power Spectral Density (PSD) of the surface, which is the Fourier transform of the autocovariance function. The integral of the PSD must be equal to the surface variance σ_s^2.

If the spatial autocovariance function of the surface height is a Gaussian, S_{HS} is also a Gaussian, thus simplifying computations. However, most optical surfaces, do *not* have a Gaussian autocovariance function. They tend to have an inverse power law falloff at large spatial frequencies. For these surfaces, the K-correlation model, or ABC model [HJBS79, CT86], is widely

used:

$$S_{HS}(f) = \frac{1}{2\sqrt{\pi}} \frac{\Gamma((c+1)/2)}{\Gamma(c/2)} \frac{ab}{\left(1+b^2f^2\right)^{\frac{c+1}{2}}}, \tag{4.17}$$

where Γ is the Gamma function. Integrating this distribution gives the surface roughness σ_s^2. It is finite only for $c > 1$ and is equal to:

$$\sigma_s^2 = \frac{\sqrt{\pi}}{c-1} \frac{\Gamma((c+1)/2)}{\Gamma(c/2)} \frac{a}{b}. \tag{4.18}$$

We use this expression to get a simpler formula for S_{HS}:

$$S_{HS}(f) = \frac{c-1}{2\pi} \frac{\sigma_s^2 b^2}{\left(1+b^2f^2\right)^{\frac{c+1}{2}}}. \tag{4.19}$$

Renormalization of $S_{\mathbf{HS}}(f)$. The scattering function S_{HS} used in Equation 4.5 is defined over the entire 2D plane for its parameter, $\boldsymbol{f}$. However, $\boldsymbol{f}$ is limited in its variation, as the projections of incoming and outgoing directions must fall in the unit disc in the plane. For energy conservation, S_{HS} must be renormalized by dividing by the integral over the set of possible values for $\boldsymbol{f}$:

$$\hat{S}_{HS} = \frac{\sigma_s^2}{\sigma_{\text{rel}}^2} S_{HS} \quad \text{with} \quad \sigma_s^2 = \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} S_{HS}(\boldsymbol{f})\,\mathrm{d}f_x\,\mathrm{d}f_y \quad \text{and} \quad \sigma_{\text{rel}}^2(\theta_i) = \int_{\|\boldsymbol{o}\|\le 1} S_{HS}(\boldsymbol{f})\,\mathrm{d}f_x\,\mathrm{d}f_y. \tag{4.20}$$

This renormalization ensures that the energy of the diffraction lobe is constant. The same renormalization constant σ_{rel} is also used for energy repartition between specular and diffraction lobes [HSCD12]:

$$A(\boldsymbol{i},\boldsymbol{o}) = \mathrm{e}^{-\left(2\pi\frac{\sigma_{\text{rel}}(\theta_i)}{\lambda}(\cos\theta_i+\cos\theta_o)\right)^2}. \tag{4.21}$$

In practice, we precompute the renormalization ratio $\sigma_{\text{rel}}/\sigma_s$ for the K-correlation model, for a large range of values for θ_i, b and c, and store it in a table, which size is 8.9 MB.

The Diffraction Lobe. To make Equation 4.10 tractable we introduce some approximations and show that the diffraction lobe is roughly a spherical convolution between the microfacet normal distribution D and the scattering function S_{HS}.

By taking into account the energy repartition $\sigma_{\text{rel}}(\theta_i)$ between the microfacet lobe and the diffraction lobe, Equation 4.10 becomes:

$$\rho_{CTD}(\boldsymbol{i},\boldsymbol{o}) = \int_\Omega \left|\frac{\boldsymbol{i}\cdot\boldsymbol{m}}{\boldsymbol{i}\cdot\boldsymbol{n}}\right| \left|\frac{\boldsymbol{o}\cdot\boldsymbol{m}}{\boldsymbol{o}\cdot\boldsymbol{n}}\right| G(\boldsymbol{i},\boldsymbol{o},\boldsymbol{m}) D(\boldsymbol{m}) (1 - A_{\sigma_{\text{rel}}}(\boldsymbol{i},\boldsymbol{o},\boldsymbol{m})) Q(\boldsymbol{i},\boldsymbol{o},\boldsymbol{m}) S_{HS}(\boldsymbol{f})\,\mathrm{d}\omega_m \tag{4.22}$$

where:

$$A_{\sigma_{\text{rel}}}(\boldsymbol{i},\boldsymbol{o},\boldsymbol{m}) = \mathrm{e}^{-\left(2\pi\frac{\sigma_{\text{rel}}(\theta_i)}{\lambda}(\boldsymbol{i}\cdot\boldsymbol{m}+\boldsymbol{o}\cdot\boldsymbol{m})\right)^2}$$

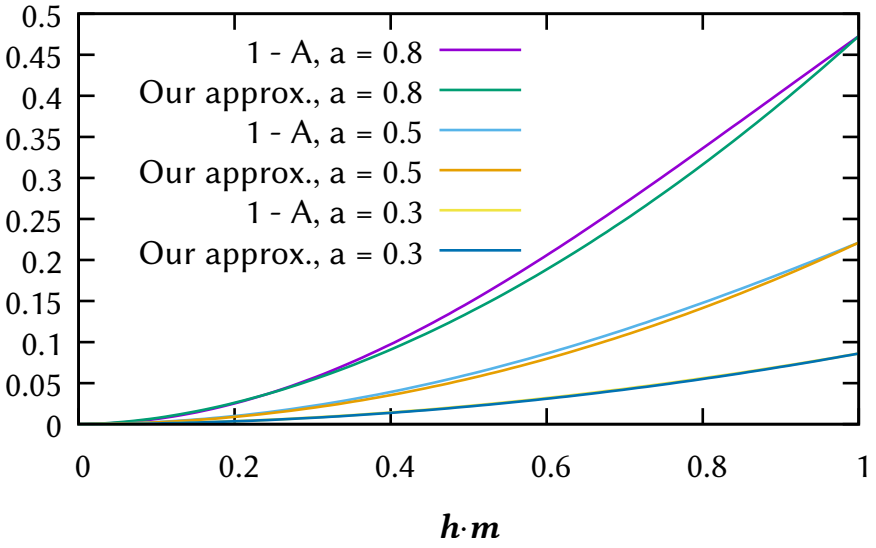


Figure 4.7: Quality of the approximation of $1 - A(\boldsymbol{i},\boldsymbol{o},\boldsymbol{m}) \approx (1 - A_{\text{spec}}(\theta_d))(\boldsymbol{h}\cdot\boldsymbol{m})^2$, depending on $a = 2\pi(2\cos\theta_d)\sigma_{\text{rel}}/\lambda$. Our approximation becomes less valid when $a > 0.5$.

We use the fact that $(\boldsymbol{i}+\boldsymbol{o}) = 2\cos\theta_d \boldsymbol{h}$, and thus for each microfacet with normal $\boldsymbol{m}$:

$$\boldsymbol{i}\cdot\boldsymbol{m} + \boldsymbol{o}\cdot\boldsymbol{m} = 2\cos\theta_d(\boldsymbol{h}\cdot\boldsymbol{m}) \tag{4.23}$$

We then approximate A with:

$$A_{\sigma_{\text{rel}}}(\boldsymbol{i},\boldsymbol{o},\boldsymbol{m}) \approx \mathrm{e}^{-\left(2\pi\frac{\sigma_{\text{rel}}(\theta_i)}{\lambda}(2\cos\theta_d)\right)^2} \tag{4.24}$$

$$A_{\sigma_{\text{rel}}(\theta_i)}(\theta_d) = \mathrm{e}^{-\left(2\pi\frac{\sigma_{\text{rel}}(\theta_i)}{\lambda}(2\cos\theta_d)\right)^2} \tag{4.25}$$

This approximation is exact for extreme values $(\boldsymbol{h}\cdot\boldsymbol{m}) = 0$ and $(\boldsymbol{h}\cdot\boldsymbol{m}) = 1$. It is good if $2\pi(2\cos\theta_d)\sigma_{\text{rel}}/\lambda < 1/2$ (see Figure 4.7).

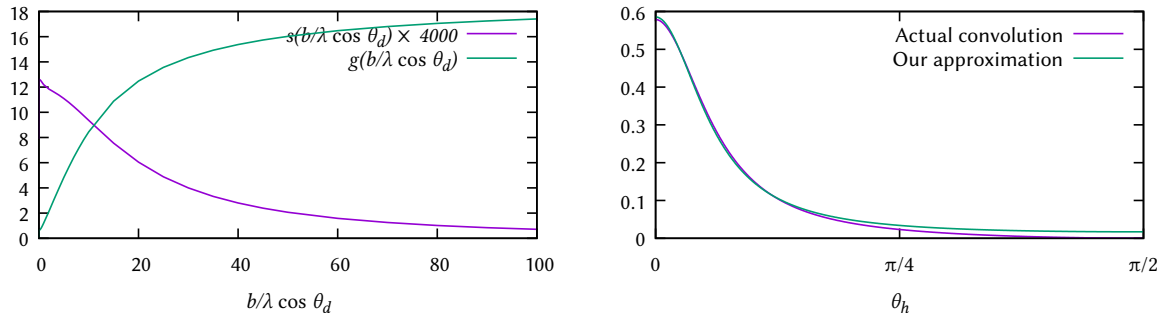
Furthermore, we also consider that the geometrical term G and the polarization factor Q can be approximated when their values are maximal (i.e., when $\boldsymbol{m} = \boldsymbol{n}$ or equivalently $f = 0$). Therefore $\rho_{CTD}(\boldsymbol{i},\boldsymbol{o})$ approximation becomes:

$$\rho_{CTD}(\boldsymbol{i},\boldsymbol{o}) \approx Q(\boldsymbol{i},\boldsymbol{o})G(\boldsymbol{i},\boldsymbol{o})(1 - A_{\sigma_{\text{rel}}(\theta_i)}(\theta_d))\int_\Omega S_{HS}(f)(\boldsymbol{h}\cdot\boldsymbol{m})^2 D(\boldsymbol{m})\,\mathrm{d}\omega_m \tag{4.26}$$

The last integral is roughly a spherical convolution between $S_{HS}(f)$ and the distribution term D, which is in our case equal to $D_{EPD}(\boldsymbol{m})$. This convolution is evaluated using Zonal Harmonics with 100 coefficients. Instead of storing 100 coefficients for each possible parameters $((\beta, p, \sigma_s, b, c))$ we approximate the convolution with a scaled version of the K-correlation model for S_{HS} :

$$(S * D)(\theta_h) \approx s(u,c,\beta,p)S_{HS}(f') \tag{4.27}$$

$$\text{with}: \quad u = \frac{b}{\lambda}\cos\theta_d \quad \text{and} \quad f' = \frac{2b}{\lambda}\sin\theta_h\, g(u,c,\beta,p). \tag{4.28}$$



(a) External and internal scaling functions s and g for $\beta = 0.03255$, $p = 2.4$, $c = 1.82$.

(b) Approximate convolution $(S * D)$ for $\theta_d = 0$, $\beta = 0.0014$, $p = 0.2$, $\sigma_s = 0.15$, $b = 1.6$, $c = 1.1$.

Figure 4.8: Evaluation the convolution $(S * D)(\theta_h)$ using s and g functions. For (a), observe that below a threshold, the functions $s(u)$ and $g(u)$ are linear: $s(u) = u$ and $g(u) = cte - u$. Consequently, the convolution (S*D) is just a scaling. Furthermore, above another threshold, g(u) is constant which implies that the convolution does not depend on θ_d nor the wavelength λ. **(b)** illustrates the fact that our approximation for the convolution is valid.

We precompute and store only the functions g and s for all parameters values (c, β, p, u) requiring a 2.5MB table. Figure 4.8(a) gives an insight of the shapes of s and g whereas Figure 4.8(b) illustrates the accuracy of our convolution approximation.

4.4 Overall Fitting Results

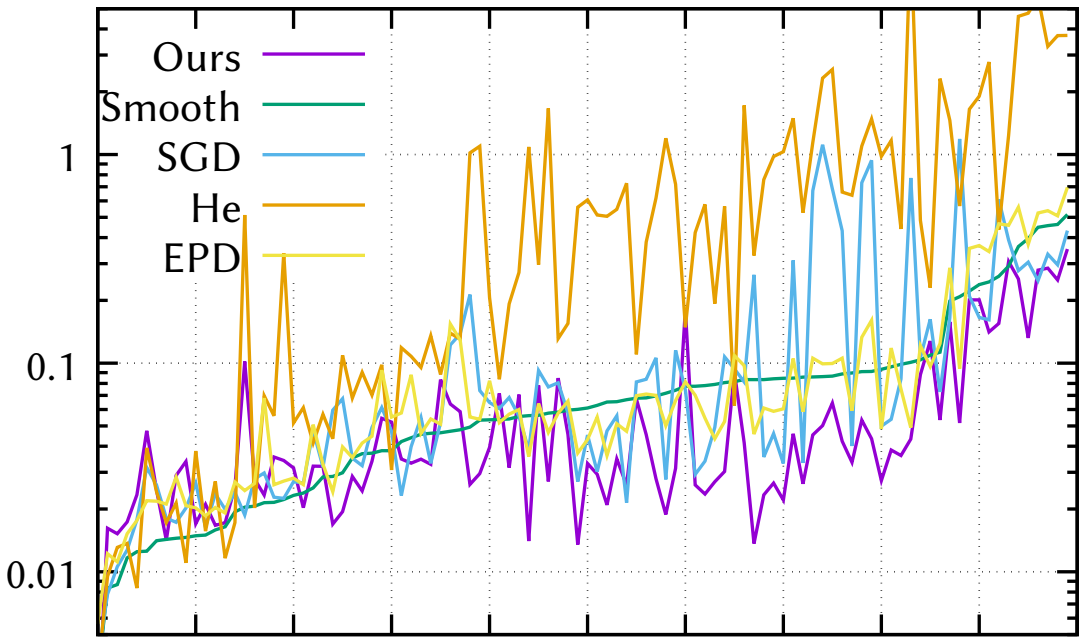


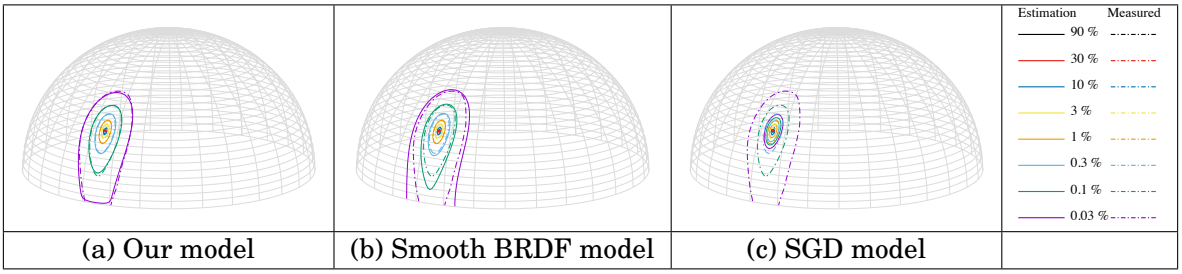
Figure 4.9: RMSE on the different tested BRDF models for all materials of the MERL database.

Figure 4.9 provides a complete summary of our fitting results. It evaluates our model and its variants (see below) for all materials of the MERL database. A detailed analysis, in BRDF and in image space, is available in the supplemental material associated with the original publication.

Looking at both qualitative (image comparison) and quantitative metrics (sMAPE, RMSE), our model provides a very good representation for glossy and specular materials. It is good, but not as much as other models for diffuse and almost diffuse materials, such as fabrics. These diffuse materials break several assumptions of the micro-facet model; taking into account multiple scattering effects would increase quality.

Specifically, our model has a smaller RMSE than other models for 62 materials, corresponding to all glossy to specular materials, from `blue-metallic-paint` to `chrome`.

4.5 Results for our conductor model



(a) Our model (b) Smooth BRDF model (c) SGD model

Figure 4.10: Lobe shapes for `blue-metallic-paint2` (blue channel) outside the incident plane, visualized as iso-lines: our model (a) provides a striking representation of BRDF behavior at all angles. The Smooth BRDF (b) model provides a good approximation but the quality decreases for grazing angles. The SGD model (c) does not have the right behavior for low values of the BRDF.

Our model excels at predicting lobe shapes for reflectance lobes. Both inside the incident plane and outside the incident plane. We represent these as isolines on the unit sphere (see Figure 4.10). As pointed by Löw et al. [LKYU12] outside the specular peak, the diffraction model provides a better prediction of lobe shape. We improve on their fitting, especially at grazing angles, by taking into account the interconnection between micro-geometry and diffraction geometry.

Figure 4.11 provides additional comparisons in image space using the sMAPE metric.

The relative intensities of diffraction effects and Cook-Torrance effects depend on the material roughness: For smooth materials (e.g., `chrome`) diffraction effects have smaller intensity than Cook-Torrance effects. The reverse is true for glossy materials (`blue-metallic-paint`, `gold-metallic- -paint`). Either way, diffraction effects play an important role in both the visual aspect and the fitting process. The EPD model has the same number of parameters as our model, with an independent diffuse lobe. Our model provides a better approximation on most materials. Lobe width can change with color channel (see Figure 4.12), in contradiction with the microfacet model. Our model explains and predicts this variation, caused by a strong diffraction lobe.

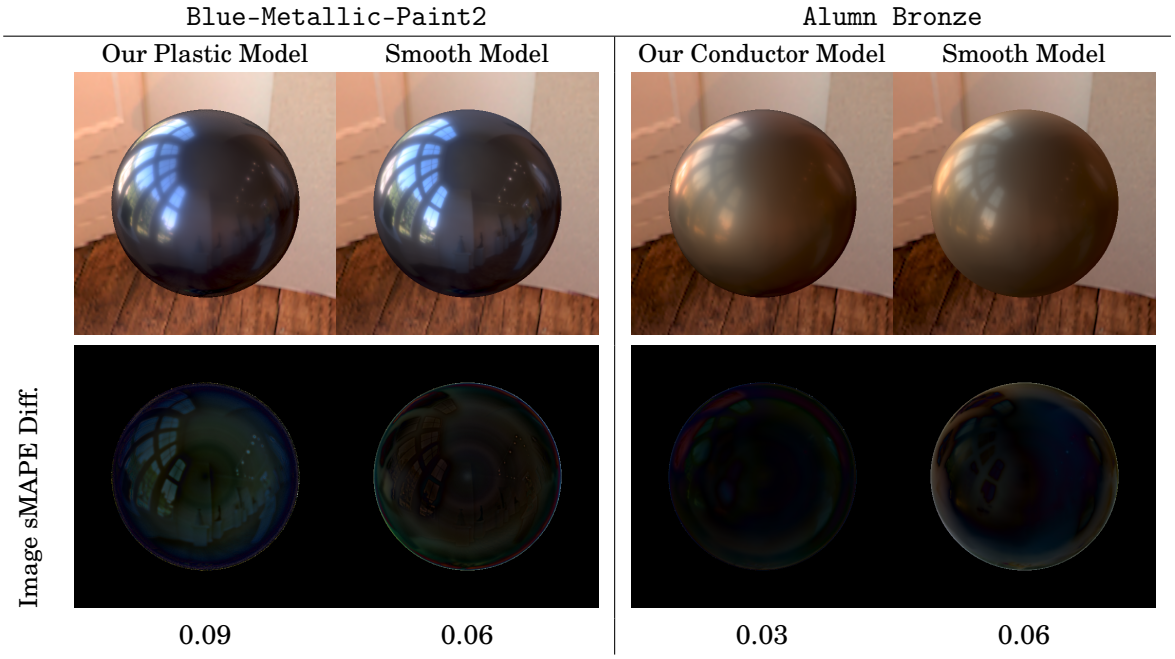


Figure 4.11: Comparisons in image space between our model and the smooth model from Löw et al. The last row shows the mean sMAPE error for the whole difference image averaged over all color channels.

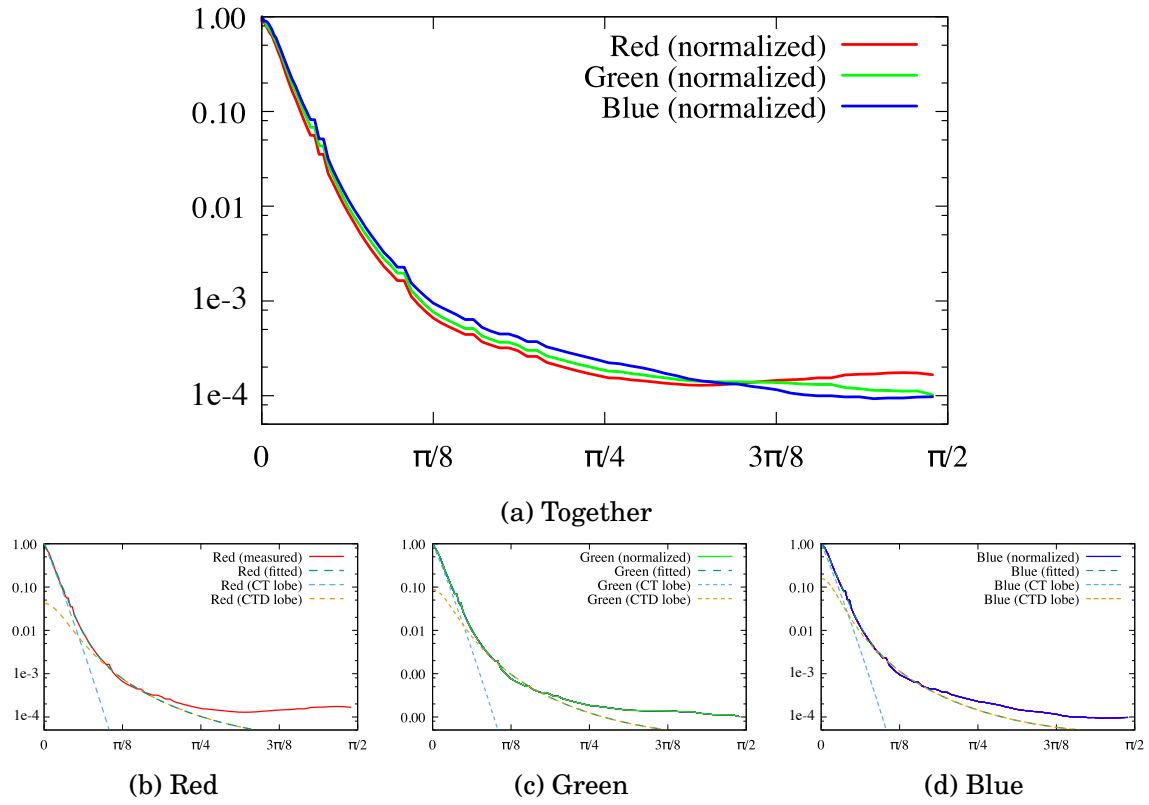
Several materials previously thought to require multiple lobes for fitting are well approximated by our model: `red-metallic-paint`, `alum-bronze`, `two-layer-gold`; one of the lobe corresponds to the microfacet lobe, the other to diffraction effects. More interestingly, both lobes have the same index of refraction.

For some materials (e.g., `gold-metallic-paint3`, `specular-or- ange-phenolic`) it is hard to separate the respective contributions of the Cook-Torrance, diffraction and diffuse components. For these, the fitting process places all contributions in one of the lobes. This gives a good visual and numerical fit, but not a good representation of the physical phenomena.

4.6 Multi-Layer Extensions

We introduce two variants (*plastic* and *subsurface*) of our combined microfacet-diffraction model to take into account materials with two layers.

The *plastic model* (see Figure 4.13 left) assumes a thin transparent layer of varnish over a diffuse layer as introduced by Weidlich and Wilkie [WW07] The *subsurface scattering model* is the BRDF approximation of subsurface scattering (see Figure 4.13 right) based on the model introduced by Jensen et al. [JMLH01]. It assumes that we have a participating media under the top interface. The top interface reflects light as in the conductor model and we add to that the light that has been refracted into the participating media and scattered multiple times before



(a) Together

(b) Red (c) Green (d) Blue

Figure 4.12: **(a)** Measured reflectance data for `nickel` as a function of θ_h, renormalized by value at origin. The width of the BRDF lobe changes with the wavelength, in contradiction with the microfacet theory. **(b)** **(c)** and **(d)**: The presence of a diffraction lobe, combined with the reflection, explains this variation of lobe width. The diffraction lobe intensity increases as the wavelength decreases (from red to blue).

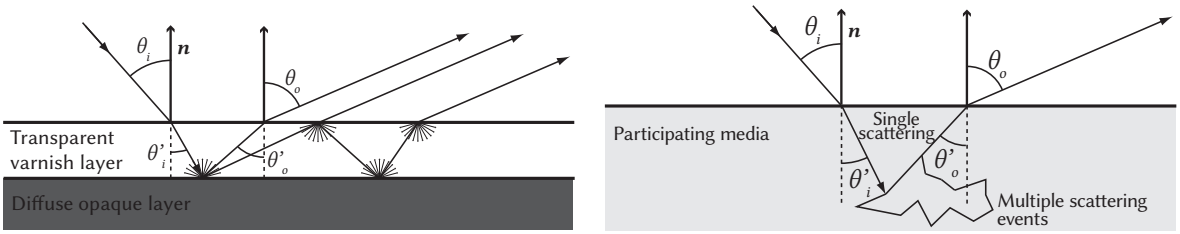


Figure 4.13: **Left**: In the *plastic model*, we assume that we have a transparent layer of refractive varnish over a diffuse material and take into account its effect with the model from Weidlich and Wilkie[WW09]. **Right**: In the *subsurface scattering model*, light is reflected on the surface as well as scattered multiple times inside the underlying participating media before leaving through the refractive interface. Images from [WW09] and

leaving. The exact formulations of these two variants are available in the supplemental material (Equations 44 to 53).

For some plastics, there is almost no difference between the measured data and what our model predicts (see Figure 4.14). It can happen that our model provides a better fit with the

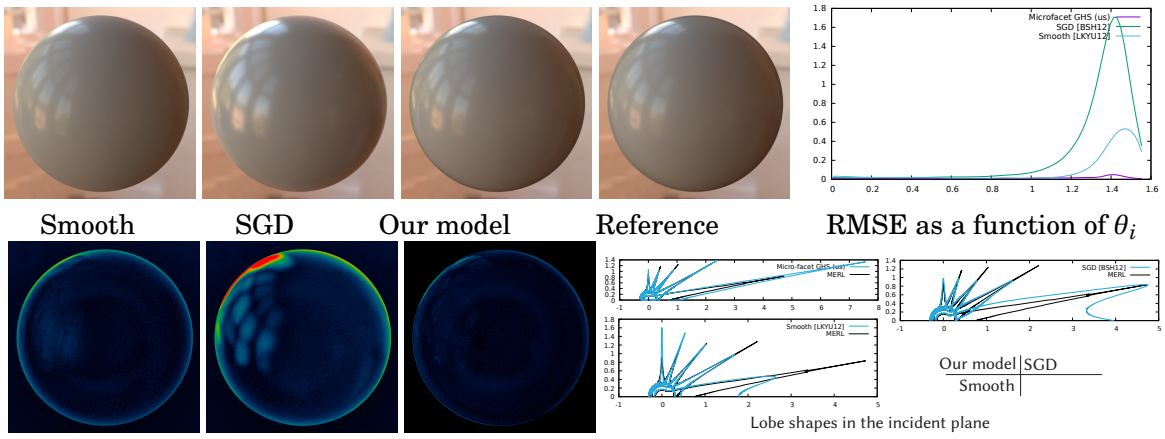


Figure 4.14: Example of a material where our model performs well: `gray-plastic`. The material behavior predicted by our model is extremely close to the measured data. Difference images using Lab color space.

measured data and lobe shape everywhere except at grazing angles (see Figure 4.15).

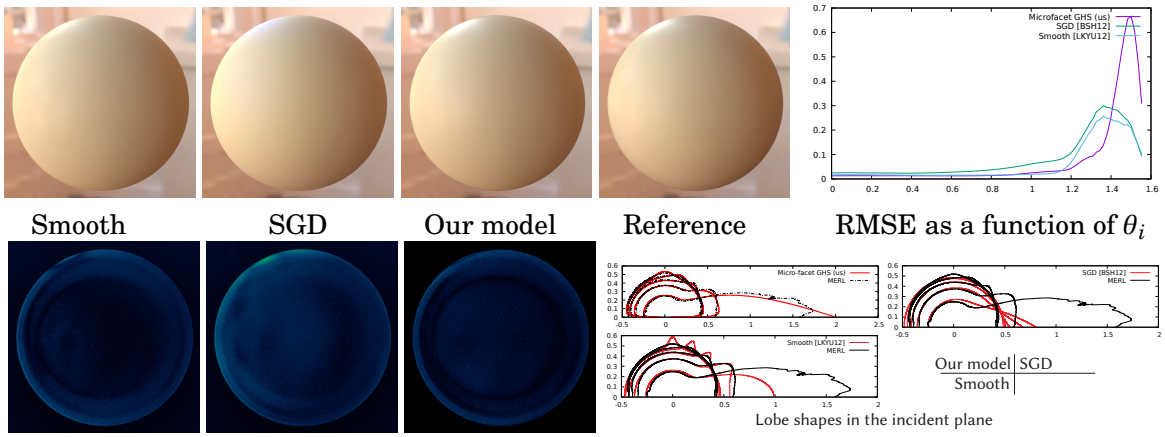


Figure 4.15: An example of a material where the RMSE for our model is larger than for others: `white-diffuse-bball`. Images and lobe shapes generated by our model are actually closer to the reference, except at grazing angles. Difference images using Lab color space.

4.7 Fitting Technique

Our fitting method combines different strategies to achieve the presented results. First, we borrows from Bagher et al. [LKYU12] by applying a log transformation and compressive weights the original measurements. This prevent the Levenberg-Marquardt numerical optimization procedure, implemented from Lourakis [Lou04], to fit higher values that are more likely to contain high error (e.g., due to the diffraction from the lens or sensor saturation). Second, we

compensate for errors at grazing angle due to the usage by Matusik et al. of a sphere as material support. Indeed, each pixel of the sensor received lights coming from a non-planar portion of the sphere. This spatial averaging translates into an angular convolution of the BRDF with the pixel footprint on the sphere. The effect of the Matusik appartus is to reduce the frequency and the maximum intensity of the measured signal (see [26] and Section 3 in the supplemental material for more details).

4.8 Comparisons with the He et al. model

He et al. [HTSG91] were probably the first to introduce a model in the Computer Graphics community that takes into account diffraction. Instead of relying on the Generalized Harvey Shack theory their model relies on the Beckmann-Kirchhoff theory. At the time, in 1991, the improved version of the Beckmann-Kirchhoff theory was not developed yet. Nevertheless, we provide a comparison with their model. However, since Ngan et al. [NDM05] did not provide their implementation nor their *exact* fitting implementation our comparison is based on our implementation (more details can be found online) and our fitting strategy (cf. previous section). We acknowledge that our fitting results are not as good as the one provided in the supplemental material from Ngan et al. We tried using their parameters (from their supplemental material), but were unable to reproduce their pictures. We suspect that scaling effects must be applied to some of the parameters.

Generally speaking, the He et al. model is accurate for "diffuse" material and can outperform our model or the one from Löw et al. [LKYU12] in terms of RMSE error in BRDF space but not necessarily in image space (see Figure 4.16 for a concrete example on the `white-fabric2` material).

4.9 Discussion and Future Work

Since 2017, the presented model has received a lot of attention (e.g., [YHW+18, CCFM22]) and some potential future work outlined here are in part carried out.

At the time of the publication the lack of spectral BRDF, as well as surface topography measurements were the main obstacles to validate *physically* our model.

This question has been partially answered very recently by Clausen et al. [CCFM22] who introduced an empirical linear shifting term that takes into account diffraction for the Cook-Torrance BRDF model. However, their BRDF spectral measurements are only done in the incident plane making their conclusion less general. Nevertheless, comparing Clausen et al. measurements with our model, without any optimization parameters, remains an interesting future work to validate further our model.

Contrary, to what is claimed in the Computer Graphics literature, the Generalized Harvey Shack theory remains accurate even for rough surfaces. Rigorous comparisons between elec-

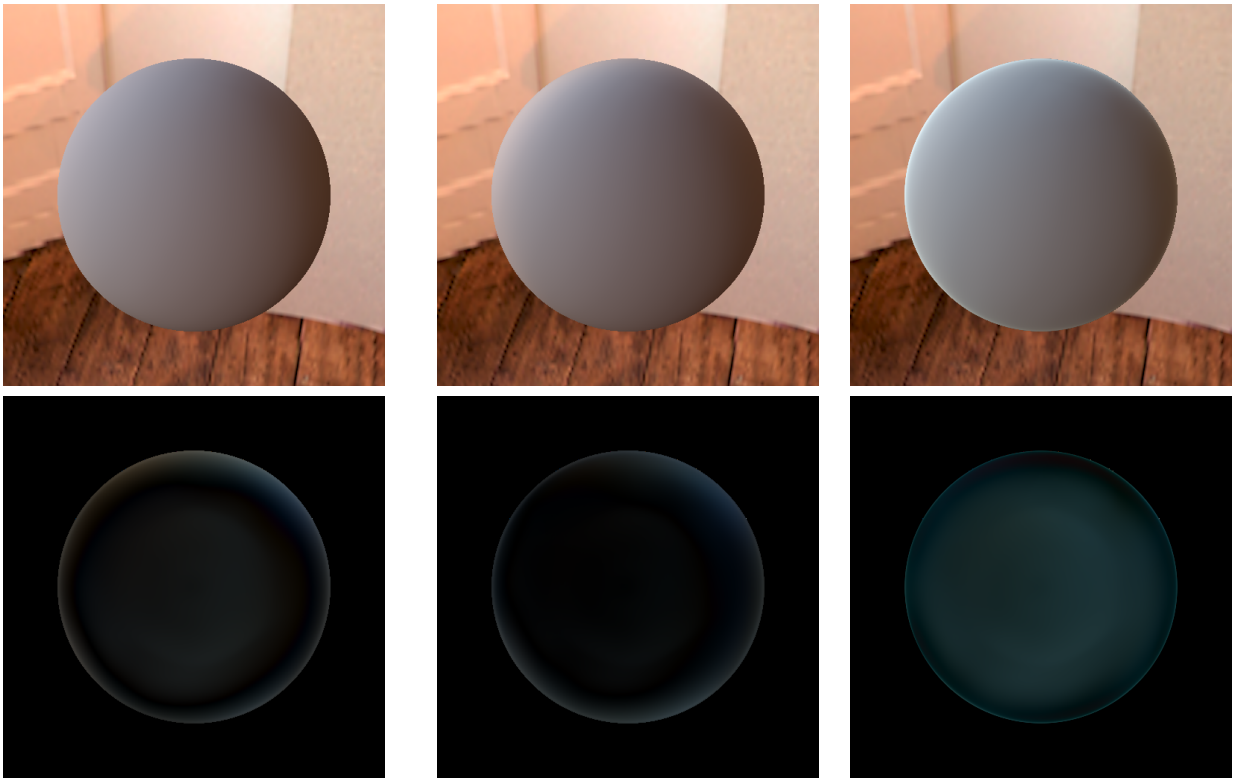


Figure 4.16: Comparisons for the `white-fabric2` material between different three diffraction models (our, Smooth from Löw et al. and He et al.) in image and in BRDF space. Although the He et al. model performs well regarding the BRDF metrics (last row) its performance in image space (RGB errors on second to last row) are below the most recent models.

tromagnetic simulations (done with Rigorous-Coupled Wave Analysis [MG81]), topography and BRDF measurements, and the GHS BRDF model have been carried out in 2011 by Schröeder et al. [SDC+11]. Moreover, despite the fact, that multiple scattering at the diffraction level is not explicitly modeled by GHS, the comparisons with RCWA tend to prove that the approximations made by GHS remain very good.

Another important issue with the Generalized Harvey Shack (GHS) theory is the polarization factor, that has been introduced empirically. Improving the diffractive GHS theory with a rigorous treatment of the polarization is an important venue for future work that requires polarized BRDF measurements. A collaboration (joint PhD supervision) with the French Onera started in Septembre 2021 on that topic.

Finally, improving our model to support anisotropy is an important and non trivial problem that has not been tackled yet. For two reasons; the first is that the convolution of the diffraction term (see Equation 4.26) with an anisotropic microfacet distribution is complex to express. The

second is the difficulty of measuring spectral and anisotropic BRDFs in order to be able to compare the developed model with the measurements.

4.A The polarization factor Q

The polarization factor was introduced by Church et al [CJZ77, CJZ79]. It depends on all angles $(\theta_i, \phi_i, \theta_o, \phi_o)$ of the BRDF as well as the index of refraction $\eta(\lambda)$. It is defined as:

$$
\begin{aligned}
Q(\eta(\lambda), \theta_i, \theta_o, \phi) &= |Q_{ss}|^2 + |Q_{sp}|^2 + |Q_{ps}|^2 + |Q_{pp}|^2 \\
Q_{ss} &= \frac{(\eta^2 - 1)\cos\phi}{(\cos\theta_i + c_i)(\cos\theta_o + c_o)} \\
Q_{sp} &= \frac{(\eta^2 - 1)c_o \sin\phi}{(\cos\theta_i + c_i)(\eta^2\cos\theta_o + c_o)} \\
Q_{ps} &= \frac{(\eta^2 - 1)c_i \sin\phi}{(\eta^2\cos\theta_i + c_i)(\cos\theta_o + c_o)} \\
Q_{pp} &= (\eta^2 - 1)\frac{c_i c_o \cos\phi - \eta^2 \sin\theta_i \sin\theta_o}{(\eta^2\cos\theta_i + c_i)(\eta^2\cos\theta_o + c_o)} \\
\text{with } \phi &= \phi_o - \phi_i - \pi, \quad c_i = \left(\eta^2 - \sin^2\theta_i\right)^{\frac{1}{2}}, \quad c_o = \left(\eta^2 - \sin^2\theta_o\right)^{\frac{1}{2}}
\end{aligned}
$$

When θ_o and θ_i are equal and when ϕ is equal to zero, Q becomes equal to the Fresnel term.

4.B Precomputed Data

The presented model requires 3 precomputed arrays:

(i) The Geometric/Shadowing term (cf. Equation 4.16) that is tabulated as 2D array of 100×1000 values, which are indexed using the parameters $(\frac{1}{\beta \tan\theta}, 5/p)$.

(ii) The renormalization factor, used for energy conservation (cf. Equation 4.21), is tabulated as 3D array of $100 \times 100 \times 100$ values, which are indexed using the parameters $(\sin\theta_i, b, c)$.

(iii) The 4D convolution (cf. Equation 4.27) is tabulated as 4D array of $25 \times 10 \times 11 \times 43$ values, which are indexed using the parameters $(p, \beta, c, \frac{b}{\lambda}\cos\theta_d)$.

All arrays are stored using double precision floating-point format.

Part III

Efficient Predictive Rendering

Generating a synthetic image requires to simulate the light transport inside a given scene, which is observed from a specific view point. In the absence of participating media, the light transport is computed by solving the rendering equation [Kaj86], which can be formulated in different forms (path, hemispherical or surface).

As stated in the introduction, my work focuses on predictive rendering. Predictive rendering consists in solving the rendering equation, for each each wavelength and each pixel of the synthetic image, without introducing any bias and ideally, also to assess the uncertainties of the computed radiance.

Despite their popularity during the 80s, finite element methods (cf. [SP94] for an historical overview) have not received much attention for the last decades due their restriction to Lambertian (perfectly diffuse) scenes and and their memory consumption.

Nowadays, most light transport simulators rely on Monte-Carlo methods (cf. [KW09] for a general introduction) to solve the rendering equation. In the context of computer generated movies, the industry has abandoned rasterization-based methods to use mostly unidirectional Path-Tracing (cf. [FHF+17]), which the simplest way to solve the rendering equation with Monte-Carlo techniques. Even the video games industry is starting to combine rasterization with ray-tracing (cf. Figure 4.2.2), which can be seen as a biased version of the path-tracing algorithms. This is certainly due to the fact that graphics card manufacturers have introduced specialized hardware (e.g., RTCore [NVi18]) units to compute visibility query very efficiently. Despite the fact that GPU compute capabilities are growing much faster than CPU ones (cf. Figure 4.2.1), generating a noise-free image with Monte-Carlo methods cannot be done interactively, even more so in real-time, without introducing some sort of bias.

Signal-processing approaches (e.g., [MMBJ17, DSHL10, HST10]), (see also [ZJL+15] for a complete overview) to denoise images have been proposed but are now becoming less popular and efficient compared to Deep Learning approaches (e.g., [CKS+17, KBS15]) based on Convolutional Neural Network (CNN). However, one must not forget two problems emanating from the later

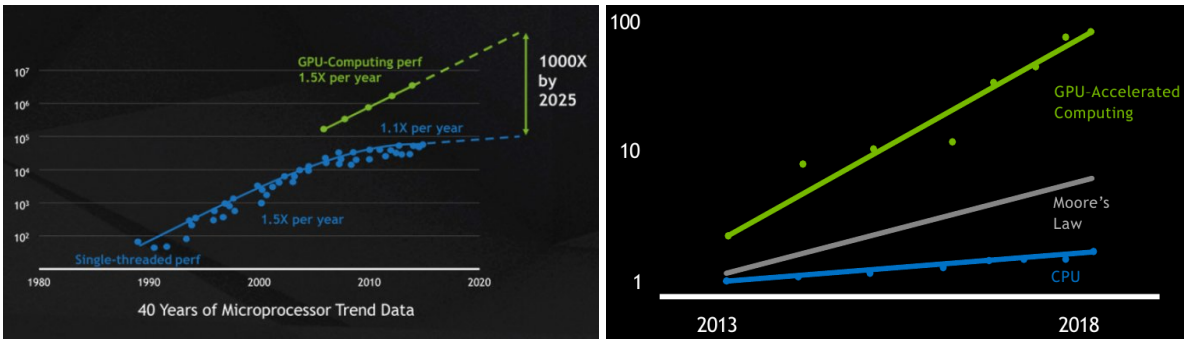


Figure 4.2.1: Graphics Processing Units (GPU) compute capabilites are growing faster than Central Processing Units (CPU), which growth is below the Moore's Law (**right**). Data from InfoWorld magazine, Martin Heller August 30, 2018.



Figure 4.2.2: . To obtain noise-free a image as above (from *Cinematic Rendering in UE4 with Real-time Ray-Tracing*) that contains multiple glossy inter-reflections in real-time, Liu et al. [LLCK19] combines a low sampling procedure (2 samples/rays for the BRDF per pixel with two-bounce reflections) with multiple denoising procedures based either on precomputed values, approximations of the rendering equation or even image-space or temporal filters. Their method shares similarity with the one introduced by Mara et al. [MMBJ17].

(i) Denoiser introduces a bias in the final image, **(ii)** In the context of predictive rendering, uncertainties, that are required, cannot be obtained from CNN-based methods. One could always argue that in the context of real-time previsualization denoising is a useful tool. To answer this question, we invite the reader to compare the images in Figure 4.2.3.

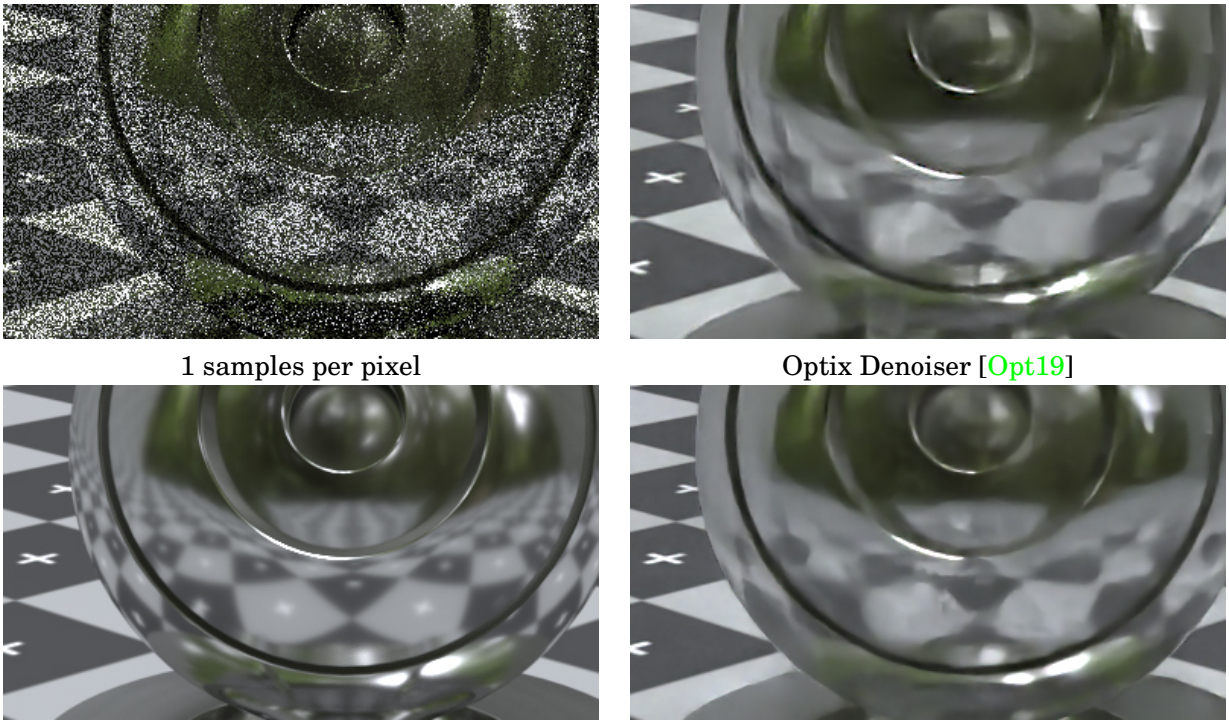
Finally, random number generation plays also an importance role in the noise reduction (e.g., [JEK+19]) and we did contribute to the field (cf. [17]) but we will not detail and discuss these types of approaches in this part. In the next section we recall how Monte-Carlo methods can be used to solve light transport without introducing any bias and how their convergence speed can be improved through importance sampling techniques.

4.3 Monte-Carlo Methods for light transport simulation

Using the Hemispherical formulation, the Rendering Equation can be stated as follows:

$$L(\boldsymbol{x} \rightarrow \boldsymbol{o}) = L_e(\boldsymbol{x} \rightarrow \boldsymbol{o}) + \underbrace{\int_{\mathscr{S}^2} f_s(\boldsymbol{x},\boldsymbol{o},\boldsymbol{\omega})(\boldsymbol{\omega}\cdot\boldsymbol{n})L(\boldsymbol{x} \leftarrow \boldsymbol{\omega})\,d\boldsymbol{\omega}}_{\text{i.e., Scattering Part}} \tag{4.29}$$

where f_s is the Bidirectional Scattering Distribution Function, which represents the directional (reflection or refraction of the light) behavior of the material at $\boldsymbol{x}$, $\mathscr{S}^2$ is the unit sphere centered



1 samples per pixel | Optix Denoiser [Opt19]

Reference | Intel Open Image Denoiser [Int19]

Figure 4.2.3: Comparisons between different types of CNN-based denoising libraries. These denoising libraries work well on average but they fail to recover high-frequency details from the glossy reflections of the material probe. Images from [Gal19].

in $\boldsymbol{x}$, $L_e(\boldsymbol{x} \rightarrow \boldsymbol{o})$ is the emitted radiance from $\boldsymbol{x}$ in the direction $\boldsymbol{o}$, $L(\boldsymbol{x} \leftarrow \boldsymbol{\omega})$ is the incoming radiance at $\boldsymbol{x}$ that arrives from direction $\boldsymbol{\omega}$ and, $L(\boldsymbol{x} \rightarrow \boldsymbol{o})$ is the radiance leaving $\boldsymbol{x}$ in the direction $\boldsymbol{o}$ that will reach the sensor. $L(\boldsymbol{x} \rightarrow \boldsymbol{o})$ needs to computed for each pixel of the image; more precisely for each pixel area.

The scattering part of the Rendering Equation:

$$I_s = \int_{\mathscr{S}^2} f_s(\boldsymbol{x}, \boldsymbol{o}, \boldsymbol{\omega})(\boldsymbol{\omega} \cdot \boldsymbol{n}) L(\boldsymbol{x} \leftarrow \boldsymbol{\omega}) d\boldsymbol{\omega} \tag{4.30}$$

is the difficult part of equation since it involves a infinite recursion to find the incident radiance. Furthermore, except for specific and well-chosen scenes, it is not possible to solve analytically the rendering equation. This comes from the visibility term (hidden in $L(\boldsymbol{x} \leftarrow \boldsymbol{\omega})$) that cannot being analytically computed. Therefore, a common choice to solve this type of equation is to rely on Monte-Carlo integration. Monte-Carlo methods are statistical tools that are used to compute high-dimensional integrals and whose only requirement is to be able to evaluate the integrand at specific position (i.e., samples).

Compared to numerical integrations methods (e.g., Rectangle, Trapezoid rule, Simpson rules, etc.) Monte-Carlo methods convergence rate (see below) is better for high dimensional integrals since it does not depend on the dimension of the integral. Although not appealing for low

dimensional problems, Monte-Carlo methods have still an advantage to manage singularities in the integrand when using importance sampling techniques.

Applying Monte-Carlo method to the scattering part of the rendering equation can be formulated as follows:

$$F_N = \frac{1}{N}\sum_{i=1}^{N}\frac{f_s(\boldsymbol{x},\boldsymbol{o},\boldsymbol{\omega}_i)(\boldsymbol{\omega}_i\cdot\boldsymbol{n})L(\boldsymbol{x}\leftarrow\boldsymbol{\omega}_i)}{\text{pdf}(\omega_i)} \tag{4.31}$$

where N is the number of samples and $\text{pdf}(\boldsymbol{\omega}_\text{i})$ is the probability density function evaluated at $\boldsymbol{\omega}_i$. F_N is called a Monte-Carlo estimator. In the following paragraphs we remind some important properties of Monte-Carlo estimators.

Error, bias, variance

The **error** of the estimator, $\text{Error}[F_N]$ is defined as the quantity:

$$\text{Error}[F_N] = F_N - I_s \tag{4.32}$$

whereas as the **bias** of the estimator is the expected value of the error:

$$\text{Bias}[F_N] = E\left[\text{Error}[F_N]\right] \quad . \tag{4.33}$$

When $\text{Bias}[F_N] = 0 \quad \forall N$, then the estimator is called **unbiased**. Equation 4.31 represents an unbiased estimator for the scattering part of the rendering equation.

$E[f(x)]$ represents the expected value of an integrable function $f(x)$ where x is a random variable:

$$E[f(x)] = \int_{-\infty}^{+\infty} f(x)\,\text{pdf}(x)\,dx \quad . \tag{4.34}$$

The **variance** of $V[f(x)]$ is defined from the expected value $E[f(x)]$:

$$\text{V}[f(x)] = E\left[\left(f(x) - E[f(x)]\right)^2\right] = E[f(x)^2] - E[f(x)]^2 \quad . \tag{4.35}$$

Observe that $\text{V}[f(x)]$ is *not* the variance of the estimator F_N, which can only be estimated with the following estimator (cf. [KW09]):

$$< V[F_N] >= \frac{1}{N-1}\left(\left(\frac{1}{N}\sum_{i=1}^{N}\left(\frac{f_s(\boldsymbol{x},\boldsymbol{o},\boldsymbol{\omega}_i)(\boldsymbol{\omega}_i\cdot\boldsymbol{n})L(\boldsymbol{x}\leftarrow\boldsymbol{\omega}_i)}{\text{pdf}(\omega_i)}\right)^2\right) - (F_N)^2\right) \tag{4.36}$$

Consistency of a Monte-Carlo Estimator

An estimator is said to be **consistent** when the error decreases toward zero with a probability of one. A sufficient condition for consistency can be formulated as follows:

$$\lim_{N\to+\infty} V[F_N] = \lim_{N\to+\infty} \text{Bias}[F_N] = 0 \tag{4.37}$$

MSE of a Monte-Carlo Estimator

It may be tempting to assume that predictive rendering requires only consistent estimators. Unfortunately this is not the case because, as stated by Veach [Vea98a], when the estimator is biased, deriving an upper bound limit of the Mean Square Error is unfeasible.

Indeed, in the general case the MSE can be written as follows:

$$\begin{aligned} MSE[F_N] &= E[(F_N - I_s)^2] = V[F_N - I_s] + (E[F_N - I_s])^2 \\ &= V[F_N] + (\text{Bias}[F_N])^2 \end{aligned}$$

In another words, to establish an upper bound of $MSE[F_N]$, one requires an upper bound on the possible bias. This is rarely possible. However, when the estimator is unbiased the $MSE[F_N]$ is directly given by the variance $V[F_N]$ or its estimate.

Convergence Rate of a Monte-Carlo Estimator

The convergence rate of Monte-Carlo sampling is slow $O(N^{-1/2})$ but, as written before, it is independent of the number of dimensions. In contrast, Simpson's rule convergence rate is $O(N^{-4/D})$ where D is the number of dimensions. When D increases the convergence rate decreases, this is known as the curse of dimensionality.

Efficiency of a Monte-Carlo Estimator

In order to compare different sampling strategies (or variance reduction techniques) it is necessary to use a metric that takes into account the variance reduction achieved *and* the required time to draw the samples; and ultimately to generate the image in the context of rendering. This metric has been introduced by Veach [Vea98b] and is called the efficiency:

$$\text{efficiency}[F_N] = \frac{1}{\text{Variance}[F_N] \cdot \text{Time}[F_N]} \quad . \tag{4.38}$$

To summarize, the Monte-Carlo theory proves that the expected value of the unbiased estimator F_N will converge toward the exact solution when N tends toward infinity:

$$E[F_N] = E\left[\sum_{i=1}^{N} \frac{f_s(\boldsymbol{x},\boldsymbol{o},\boldsymbol{\omega}_i)(\boldsymbol{\omega}_i \cdot \boldsymbol{n}) L(\boldsymbol{x} \leftarrow \boldsymbol{\omega}_i)}{\text{pdf}(\omega_i)}\right] = \int_{\mathscr{S}^2} f_s(\boldsymbol{x},\boldsymbol{o},\boldsymbol{\omega})(\boldsymbol{\omega} \cdot \boldsymbol{n}) L(\boldsymbol{x} \leftarrow \boldsymbol{\omega}) d\boldsymbol{\omega} \tag{4.39}$$

Improving Efficiency: Importance Sampling Techniques

The choice of the pdf has a great impact on the efficiency of the estimator. Ideally, we would like to draw samples proportionally to the integrand (i.e., the numerator of Equation 4.31). Unfortunately, in the generic case, this is impossible because the incoming radiance is unknown; it is precisely the quantity we would like to compute (i.e., we are still stuck with a recursion problem).

However, the material BSDFs are known before solving the Rendering Equation,this means that drawing samples according to the BSDF is achievable (i.e., the pdf is proportional to

$f_s(\boldsymbol{x},\boldsymbol{o},\boldsymbol{\omega}_i)$). This is called **BSDF importance sampling**. Depending on the form of the BSDF one can rely on analytical and closed-form expressions to draw samples from the BSDF (or even $f_s(\boldsymbol{x},\boldsymbol{o},\boldsymbol{\omega}_i)\cdot(\boldsymbol{\omega}_i\cdot\boldsymbol{n})$). When no closed-form expression exists one can rely on tabulating (e.g., [LRR04]) the probability density function (pdf) and the cumulative density function ()cdf). The rational framework presented in Chapter 3 has been used with success to represent and compress tabulated pdf and inverse cdf. The framework is generic and can be used to compressed any kind of inverse cdf or pdf regardless for the complexity of the product $f_s(\boldsymbol{x},\boldsymbol{o},\boldsymbol{\omega}_i)\cdot(\boldsymbol{\omega}_i\cdot\boldsymbol{n})$.

In the case where some pre-computations (e.g., Photon-Mapping [Jen01]) are possible, one gain access to information about the incoming light and can therefore exploit this knowledge to draw samples from $L(\boldsymbol{x}\leftarrow\boldsymbol{\omega})$, this is called **Light Source importance sampling**.

As we can see, it is difficult to generate samples following a probability law proportional to the integrand of the rendering equation, we can naturally ask the question if it is possible to combine samples generated according to the different part (BSDF or light sources). This is the interest and the motivation of **Multiple Importance Sampling (MIS)**.

Multiple importance sampling has been studied and investigated a lot since its introduction [Vea98a]. The main idea behind MIS is to combine (without introducing any bias) samples from different importance sampling techniques. This works especially well in the context of direct lighting where the light source positions are known. One can therefore combine samples from the light source with samples from the BSDF. The next chapter presents our contribution on how to to improve MIS to balance the number of samples between the light source and the BSDF more efficiently.

Finally, there are other importance sampling techniques that try to generate samples *approximately* to a desired pdf. These techniques (e.g., ReSTIR [BWP+20], ReSTIR GI [OLK+21]) have regain interest especially the ones based on Resampled Importance Sampling (cf. [TCE05]). The main idea behind Resampled Importance Sampling is to use a *proxy* function which pdf is known to sample a target function, which pdf is unknown. Despite their impressive efficiency, ReSTIR is limited to direct lighting and ReSTIR GI remains biased. Therefore, they cannot be used in the context of predictive rendering. A notable exception is the recent publication of Lin et al. [LKB+22] where they generalized resampling importance sampling (GRIS). Their main contribution is to introduce a correct way to compute the weights when re-using samples between pixels of the rendered image. This type of work is orthogonal to ours, and although not trivial, the technique presented in the next chapter, could be combined with GRIS.

CHAPTER 5

PER-PIXEL OPTIMIZED BALANCING FOR MULTIPLE IMPORTANCE SAMPLING

5.1 Motivation

One of the popular choice to compute efficiently the rendering equation is to combined different importance sampling strategies by using Multiple Importance Sampling [Vea98a]. Given S sampling strategies, an MIS estimator F_N to compute the reflective part of the rendering equation:

$$L(\boldsymbol{x} \rightarrow \boldsymbol{o}) = \int_{\mathscr{S}^2} \rho(\boldsymbol{x},\boldsymbol{o},\boldsymbol{\omega})(\boldsymbol{\omega}\cdot\boldsymbol{n})L(\boldsymbol{x} \leftarrow \boldsymbol{\omega})\,d\boldsymbol{\omega} \tag{5.1}$$

is defined by:

$$L(\boldsymbol{x} \rightarrow \boldsymbol{o}) \approx F_N = \sum_{s=1}^{S} \frac{1}{N_s} \sum_{i=1}^{N_s} \mathrm{w}_s(\boldsymbol{\omega}_{i,s}) \frac{\mathrm{f}(\boldsymbol{x},\boldsymbol{o},\boldsymbol{\omega}_{i,s})}{\mathrm{p}_s(\boldsymbol{\omega}_{i,s})} \tag{5.2}$$

where $\mathrm{f}(\boldsymbol{x},\boldsymbol{o},\boldsymbol{\omega}) = \rho(\boldsymbol{x},\boldsymbol{o},\boldsymbol{\omega})(\boldsymbol{\omega}\cdot\boldsymbol{n})L(\boldsymbol{x} \leftarrow \boldsymbol{\omega})$, p_s represents the probability density function associated with the s-th sampling strategy, $N = \sum_s N_s$, and w_s are weights based on different heuristic (e.g., max, power, balanced). By exploiting different sampling strategies, the MIS estimator relies on the hope that at least one of them performs well, for a limited overhead as it should not perform worse than any of them.

For the balance heuristic, the weights w_s are given by:

$$\mathrm{w}_s(\boldsymbol{\omega}_{i,s}) = \frac{N_s\,\mathrm{p}_s(\boldsymbol{\omega}_{i,s})}{\sum_{s=1}^{N_s} N_s \mathrm{p}_s(\boldsymbol{\omega}_{i,s})}. \tag{5.3}$$

For the special (but common) case where only two sampling strategies (e.g., BRDF and light source sampling) are available, the weights for the balance heuristic become:

$$\mathrm{w}_s(\boldsymbol{\omega}_{i,s}) = \frac{N_s\,\mathrm{p}_s(\boldsymbol{\omega}_{i,s})}{N_{\mathrm{brdf}}\mathrm{p}_{\mathrm{brdf}}(\boldsymbol{\omega}_{i,\mathrm{brdf}}) + N_{\mathrm{light}}\mathrm{p}_{\mathrm{light}}(\boldsymbol{\omega}_{i,\mathrm{light}})} \tag{5.4}$$

where N_{brdf} (resp. N_{light}) represents the number of samples allocated for BRDF (resp. light source) importance sampling and, the MIS estimator F_N can be written as:

$$F_N = F_{N_{\text{brdf}},N_{\text{light}}} = \frac{1}{N}\sum_{i=1}^{N} \frac{f(\boldsymbol{x},\boldsymbol{o},\boldsymbol{\omega}_{i,s})}{N_{\text{brdf}}\text{p}_{\text{brdf}}(\boldsymbol{\omega}_{i,\text{brdf}}) + N_{\text{light}}\text{p}_{\text{light}}(\boldsymbol{\omega}_{i,\text{light}})} \tag{5.5}$$

with :

- N, the total number of samples, is equal to $N = N_{\text{brdf}} + N_{\text{light}}$
- $\text{p}_{\text{brdf}}(\boldsymbol{\omega}_{i,\text{brdf}})$ represents the probability density function to sample the direction $\boldsymbol{\omega}_{i,\text{brdf}}$, obtained by sampling the BRDF at the position $\boldsymbol{x}$.
- $\text{p}_{\text{light}}(\boldsymbol{\omega}_{i,\text{light}})$ represents the probability density function to sample the direction $\boldsymbol{\omega}_{i,\text{light}}$, which is obtained by sampling a position on a light source.

Since $N_{\text{brdf}} \leqslant N$ and represents a fraction of the total number of samples, as pointed out by Georgiev [GKPS12], the estimator of equation 5.5 can be rewritten as:

$$F_N(\alpha) = \sum_{i=1}^{N} \frac{\text{f}(\boldsymbol{x},\boldsymbol{o},\boldsymbol{\omega}_{i,s})}{\alpha\text{p}_{\text{brdf}}(\boldsymbol{\omega}_{i,\text{brdf}}) + (1-\alpha)\text{p}_{\text{light}}(\boldsymbol{\omega}_{i,\text{light}})} \quad \text{with}\, \alpha \in [0,1]. \tag{5.6}$$

This is called a Defensive Importance Sampling (DIS) strategy [Hes95]. The most important observation here, is that a set of N samples (i.e., $\boldsymbol{\omega}$ directions/paths) are generated according to a probability density function $p(\boldsymbol{\omega}_{i,s})$ equal to

$$p_\alpha(\boldsymbol{\omega}_{i,s}) = \alpha\,\text{p}_{\text{brdf}}(\boldsymbol{\omega}_{i,s}) + (1-\alpha)\,\text{p}_{\text{light}}(\boldsymbol{\omega}_{i,s}). \tag{5.7}$$

Moreover, α represents the sample ratio $\frac{N_{\text{brdf}}}{N}$ whereas $(1-\alpha)$ represents the sample ratio $\frac{N_{\text{light}}}{N}$. In another words, specifying α (resp. $(1-\alpha)$) is equivalent to define the number of samples used to sample the BRDF (resp. the light sources).

Although Veach [Vea97] demonstrated that in *average* the balance heuristic (i.e., $\alpha = 0.5$) gives the best results, there is no reason that locally, another balancing strategy (i.e., another value of α) could reduce further the variance.

Our contribution [15] is to compute a value for $\alpha(x)$ per pixel (or 3D position $\boldsymbol{x}$) that minimizes the variance (i.e., the second order) of the Defensive Estimator of Equation 5.6.

In the next section we explain how to estimate α with a subset of samples (called learning samples) and we introduce a new Monte-Carlo estimator to compute the final value of the radiance $L(x \rightarrow \boldsymbol{o})$. Sections 5.3 through 5.5 provide detailed explanations and implementation. In section 5.6, we show that for our test scenes and for direct lighting, our approximation also improves the average variance, validating its accuracy. We also demonstrate that we can use our approach with a better efficiency compared to previous balancing techniques. After a discussion (Section 5.7), we introduce recent improvements (Section 5.8) and conclude this chapter with some short term future work (Section 5.9).

5.2 Second-order Variance Reduction

The variance $\mathrm{V}[F_N(\alpha)]$ of the previous estimator $F_N(\alpha)$, can be reformulated with respect to the variance of the sampling strategy (i.e., $\mathrm{V}[F_1(\alpha)]$.

Let $f_i(\alpha) = \frac{\mathrm{f}(\boldsymbol{x},\boldsymbol{o},\boldsymbol{\omega}_{i,s})}{p_\alpha(\boldsymbol{\omega}_{i,s})}$, then

$$\begin{aligned}
V[F_N(\alpha)] &= V\left[\frac{1}{N}\sum_{i=1}^{N} f_i(\alpha)\right] \\
&= \frac{1}{N^2} V\left[\sum_{i=1}^{N} f_i(\alpha)\right] \\
&= \frac{1}{N^2}\sum_{i=1}^{N} V[f_i(\alpha)] \\
&= \frac{1}{N^2}\underbrace{\left(V[f_1(\alpha)] + V[f_2(\alpha)] + \ldots + V[f_N(\alpha)]\right)}_{NV[f_1(\alpha)]} \quad \text{(Variance sum of } N \text{ independent random variables)}
\end{aligned}$$

Hence :

$$V[F_N(\alpha)] = \frac{1}{N} V[f_1(\alpha)] = \frac{1}{N} V[F_1(\alpha)] \tag{5.8}$$

This means that finding the best α that minimizes the variance $V[F_N(\alpha)]$ is equivalent to minimize the sampling variance of the strategy $V[F_1(\alpha)]$. Furthermore, we observe that when different sampling strategies (e.g., p_{brdf} and p_{light}) are unbiased the weighted combined strategy $p_\alpha(\boldsymbol{\omega}_{i,s})$ is also unbiased ($E[F_1(\alpha)] = L(\boldsymbol{x} \rightarrow \boldsymbol{o})$) and is therefore independent of the choice of α.

Since $V[F_1(\alpha)]$ is equal to:

$$V[F_1(\alpha)] = E[F_1(\alpha)^2] - E[F_1(\alpha)]^2 \tag{5.9}$$

to minimize the variance of one sample we need to minimize the expected value:

$$\mathrm{Min}_\alpha(E[F_1(\alpha)^2]) \tag{5.10}$$

$$\text{with} \quad E[F_1(\alpha)^2] = \int \left(\frac{f(\boldsymbol{x},\boldsymbol{o},\boldsymbol{\omega}_{1,s})}{p_\alpha(\boldsymbol{\omega}_{1,s})}\right)^2 p_\alpha(\boldsymbol{\omega}_{1,s})\, d\boldsymbol{\omega} = \int \frac{f(\boldsymbol{x},\boldsymbol{o},\boldsymbol{\omega}_{1,s})^2}{p_\alpha(\boldsymbol{\omega}_{1,s})}\, d\boldsymbol{\omega} \tag{5.11}$$

However, in the rendering context no analytical form is available to solve the problem in Equation 5.10 and we therefore compute an approximated solution using a second-order Taylor expansion around $\alpha = 1/2$. The choice of Taylor expansion around $\alpha = 1/2$ is motivated by the following reason. It ensures that the best approximation is centered in the definition domain of α, which is $[0,1]$. Similarly, an approximation around $\alpha = 0$ or $\alpha = 1$ will give a strong priority to one particular sampling technique

To solve the optimization problem we need to find for which value of α the gradient of $E[F_1(\alpha)^2]$) wrt. to α vanishes:

$$\nabla_\alpha\left(E\left[(F_1(\alpha))^2\right]\right) = 0$$

We demonstrate in Section 1.3 of the supplemental material of [15]) that this is achieved for:

$$\begin{aligned} \alpha &= \frac{1}{2} + \frac{I_1}{4\,I_2} \\ I_1 &= 2\int \frac{\left(\mathrm{p}_{\mathrm{brdf}}(\boldsymbol{\omega}_i) - \mathrm{p}_{\mathrm{light}}(\boldsymbol{\omega}_i)\right)}{\left(\mathrm{p}_{\mathrm{brdf}}(\boldsymbol{\omega}_i) + \mathrm{p}_{\mathrm{light}}(\boldsymbol{\omega}_i)\right)^2} f^2(\boldsymbol{x},\boldsymbol{o},\boldsymbol{\omega}_i)\, d\boldsymbol{\omega}_i \\ I_2 &= 2\int \frac{\left(\mathrm{p}_{\mathrm{brdf}}(\boldsymbol{\omega}_i) - \mathrm{p}_{\mathrm{light}}(\boldsymbol{\omega}_i)\right)^2}{\left(\mathrm{p}_{\mathrm{brdf}}(\boldsymbol{\omega}_i) + \mathrm{p}_{\mathrm{light}}(\boldsymbol{\omega}_i)\right)^3} f^2(\boldsymbol{x},\boldsymbol{o},\boldsymbol{\omega}_i)\, d\boldsymbol{\omega}_i \end{aligned} \tag{5.12}$$

Computing α requires to evaluate two integrals I_1 and I_2 that can only be computed by generating samples (path) since $f(\boldsymbol{x},\boldsymbol{o},\boldsymbol{\omega}_{i,s})$ is the triple product of the reflectance times lighting times cosine factor. I_1 and I_2 are evaluated with M samples by using a classical Monte-Carlo estimator:

$$F_{\frac{M}{2},\frac{M}{2}}[\alpha] = \frac{1}{4}\left(2 + \frac{F_{\frac{M}{2},\frac{M}{2}}[I_1]}{F_{\frac{M}{2},\frac{M}{2}}[I_2]}\right) \tag{5.13}$$

$$F_{\frac{M}{2},\frac{M}{2}}[I_1] = \frac{1}{M}\sum_{i=1}^{M} \frac{\left(\mathrm{p}_{\mathrm{brdf}}(\boldsymbol{\omega}_i) - \mathrm{p}_{\mathrm{light}}(\boldsymbol{\omega}_i)\right)}{\left(\mathrm{p}_{\mathrm{brdf}}(\boldsymbol{\omega}_i) + \mathrm{p}_{\mathrm{light}}(\boldsymbol{\omega}_i)\right)^2} \frac{f^2(\boldsymbol{x},\boldsymbol{o},\boldsymbol{\omega}_i)}{\left(\frac{M}{2}\,\mathrm{p}_{\mathrm{brdf}}(\boldsymbol{\omega}_{i,\mathrm{brdf}}) + \frac{M}{2}\,\mathrm{p}_{\mathrm{light}}(\boldsymbol{\omega}_{i,\mathrm{light}})\right)} \tag{5.14}$$

$$F_{\frac{M}{2},\frac{M}{2}}[I_2] = \frac{1}{M}\sum_{i=1}^{M} \frac{\left(\mathrm{p}_{\mathrm{brdf}}(\boldsymbol{\omega}_i) - \mathrm{p}_{\mathrm{light}}(\boldsymbol{\omega}_i)\right)^2}{\left(\mathrm{p}_{\mathrm{brdf}}(\boldsymbol{\omega}_i) + \mathrm{p}_{\mathrm{light}}(\boldsymbol{\omega}_i)\right)^3} \frac{f^2(\boldsymbol{x},\boldsymbol{o},\boldsymbol{\omega}_i)}{\left(\frac{M}{2}\,\mathrm{p}_{\mathrm{brdf}}(\boldsymbol{\omega}_{i,\mathrm{brdf}}) + \frac{M}{2}\,\mathrm{p}_{\mathrm{light}}(\boldsymbol{\omega}_{i,\mathrm{light}})\right)} \tag{5.15}$$

These samples, named learning samples, should not be wasted. One solution would be to store and use them with the estimated α to compute the shading. However, this would require an extra-storage of size M per pixel.

Therefore, we prefer to use a less optimal but simpler solution that reduces the memory footprint. We use M samples to approximate α (by evaluating in parallel Equations 5.12) and reuse these evaluations for our final estimate of the radiance. This is achieved by introducing the following Monte-Carlo estimator :

$$F_N^{\mathrm{final}}\left[L(\boldsymbol{x}\to\boldsymbol{o})\right] = \frac{M}{N} F_{\frac{M}{2},\frac{M}{2}}\left[L(\boldsymbol{x}\to\boldsymbol{o})\right] + \frac{N-M}{N} F_{\alpha N,(1-\alpha)N}\left[L(\boldsymbol{x}\to\boldsymbol{o})\right] \tag{5.16}$$

- $F_{\frac{M}{2},\frac{M}{2}}$ represents an estimator used to approximate α with the balance heuristic (cf. Equation 5.13). The number of samples assigned for BRDF (resp. light source) sampling strategy is set to $M/2$.

- $F_{\alpha N,(1-\alpha)N}$ represents the DIS estimator where the number of samples to evaluate the BRDF (resp. light source) contribution is proportional to αN (resp. $(1-\alpha)\cdot N$).

We make two additional remarks:

- N is the total number of samples used to estimate the radiance including the M learning samples.

- Therefore, the efficiency of the final estimator F_N^{final}, averaged over all pixels of the image, is the most important statistic to compare our technique with the previous ones such as the balance, max or power heuristics.

5.3 Theoretical analysis

If we consider the extreme case where the sampling strategy should be set to full BRDF sampling in the case of a Dirac reflectance function, replacing $p_{\text{light}} = 0$ in Equation 5.12 yields $\alpha = \frac{3}{4}$ instead of the theoretical value $\alpha = 1$.

Similarly, consider the extreme case where the sampling strategy should be set to full light source sampling only in the case of a Dirac emitter, by replacing $p_{\text{brdf}} = 0$ in Equation 5.12 we obtain $\alpha = \frac{1}{4}$ instead of the theoretical value $\alpha = 0$.

Therefore, we know that for these two extreme cases the optimal value for α will not be found. However, classically for MIS, it is usual to avoid using either $\alpha = 0$ or $\alpha = 1$ to ensure that the resulting estimator remains unbiased even if one of the two techniques is biased. Furthermore, the value of α still clearly indicates the dominant strategy.

One can show (cf. Section 1.5 in the supplemental material of [15]) that our variance reduction technique is in fact proportional to $|2\alpha - 1|$. The closer the estimated α is to 1/2, the less efficient becomes our approach.

5.4 Convergence speed for α computation

Equation 5.16 raises the question about the number of samples M required to obtain a "good" approximation for α. Ideally we would like to keep M low to avoid penalizing the overall efficiency of our technique. As shown in Figure 5.5.1 we found empirically that for 11 tested scenes that $M = 128$ is a reasonable value.

5.5 α clamping

As discuss above, we know that our estimation is only an approximation of the best α. In order to investigate its limited accuracy in a more general context, we have tested (cf. Figure 5.5.2) the effect of clamping the estimated α. Empirically, Figure 5.5.2 shows that for some scenes, the variance increases if the value of α is outside the interval [0.025,0.975] (i.e., with a size of 0.95). Furthermore, this experience confirms two facts. First, the second-order approximation is not sufficient to capture some variance oscillations in the interval [0,1] since clamping improves the results. Second, it also confirms that our estimation still provides a good hint of what is the

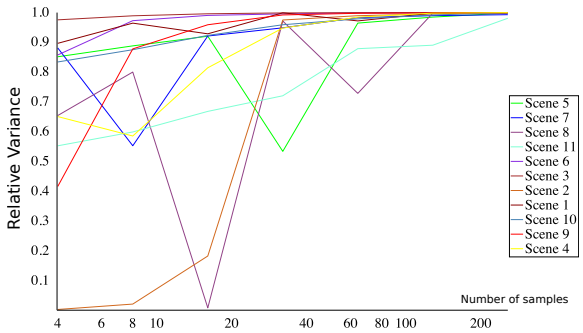


Figure 5.5.1: Number of samples M to estimate α per scene. The y-axis represents the variance of $F_{N=256,\alpha}$, relatively to the best variance obtained with our technique for different value of M for each scene, thus illustrating the convergence of α estimation.

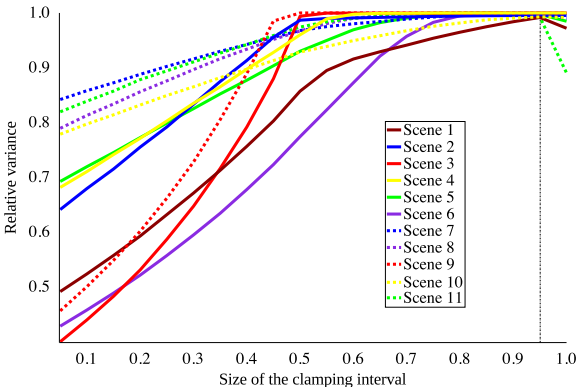


Figure 5.5.2: Effect of clamping α for each scene. The y-axis represents the variance for $F_N(\alpha)$, relatively to the best variance. The x-axis represents the size of an interval centered at 1/2 that restricts the value of α. For each scene, the number of samples per light source to estimate α is set to $M = 128$ and the one to compute the variance to $N = 256$.

dominant strategy. Finally, it justifies the clamping of α to [0.025,0.975] for all the following results.

5.6 Results

All results of this section were produced by modifying the direct illumination integrator in the second version of PBRT [PH10]. All the results are computed using 4 camera rays per pixel. For each camera ray, we use the default PBRT approach and sample uniformly all light sources and accumulate the contribution of each light source to estimate the final radiance. We tested our method on 11 scenes (cf. Figures 5.6.1 and 5.6.2) that exhibit very different behaviours in terms of light transport complexity.

Table 5.6.1 presents different metrics computed for each scene and for different sampling strategies. We do not compare with the α-max heuristic [GKPS12] since it requires additional knowledge about which strategy achieves more often a better sampling than the others.

As shown at the bottom of the table our method always outperform the balance, the maximum or the power heuristics. To complete our efficiency analysis we also compare our method when using only one sampling strategy: either BRDF or light source sampling.

For scenes with mostly diffuse BRDFs (Scenes 2,3,4) using only the light sampling strategy improves significantly the efficiency. Similarly, for a scene, such as scene 9, strongly dominated by glossy BRDF under low-frequency lighting, the BRDF sampling strategy is more efficient than our method. However, choosing only either the light or BRDF sampling strategy requires

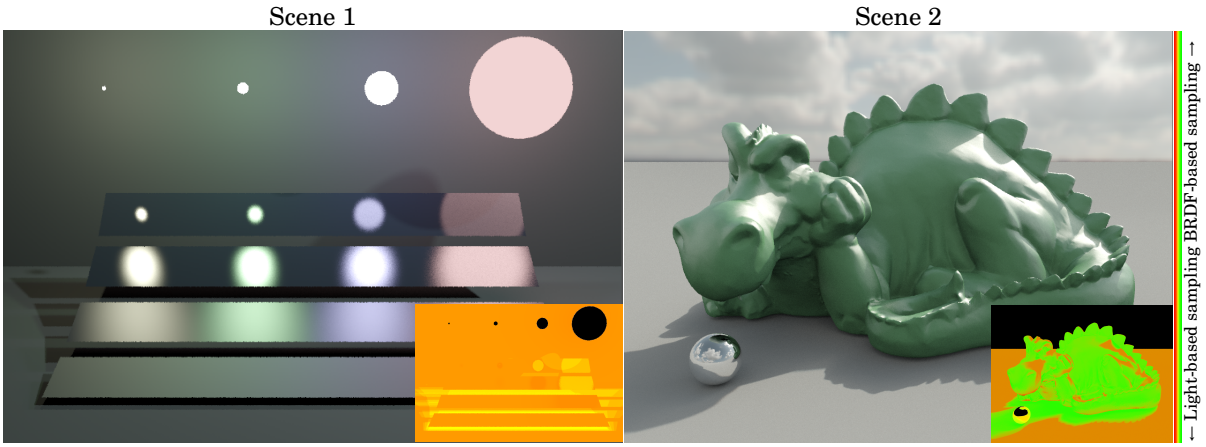


Figure 5.6.1: Our per-pixel second-order approximation of the variance leads to a new and automatic approach for balancing the number of samples between two different sampling strategies. Except for light sources, the inset images show the sample distribution for each pixel. Yellow corresponds to the default balance heuristic strategy [Vea97]. Compared to the balance heuristic, the variance is reduced by (**Left**) 26% and (**Right**) 20% in average (14% and 11% for the standard deviation).

before rendering the image requires additional precomputed knowledge that are not available in practice.

Finally, note that all numbers present in the aforementioned table are average values computed over the whole image. Figure 5.6.3 shows more precisely where our method outperforms (red pixels) the balance heuristic, where it does not (green pixels) and where they are equal (black pixels).

5.7 Discussion

Pajot et al. [PBPP11] have introduced the empirical notion of representativity that estimates the concentration information for BRDF or photon maps. Intuitively, narrow lobe-shaped PDFs will request more samples. General guidelines to develop new representativity functions are introduced, but problems may occur when multiple concentrations exist such as in environment maps. Therefore, our work differs on the following points. First, we do not need to introduce a dedicated representativity function when introducing a new sampling strategy. Second, we do not make any assumption on the shape of the PDFs, which is difficult to achieve on generic environment maps. Third, our balancing strategy is completely done per-pixel and takes into account the local shading configuration.

Since our contribution [15], many sampling strategies (e.g. [KVG+19, GGSK19, HEV+16a]) have been proposed, along with solutions to determine which one is good. However, most of these solutions focus on finding efficient strategies using equal sample distribution and only Sbert et

	Method \ Scene	1	2	3	4	5	6	7	8	9	10	11
Variance	BRDF	74.8	966	251	10.9	1.14	10.5	0.94	48.1	**0.08**	347	6.87
	Light	0.52	**0.54**	**0.05**	**0.52**	10.1	0.73	10.7	0.84	472	1.93	2.31
	Balanced heuristic	1	1	1	1	1	1	1	1	1	1	1
	Power heuristic	1.10	1.02	1.01	1.11	1.10	1.15	1.11	1.07	1.02	1.09	1.07
	Maximum heuristic	1.38	1.07	1.01	1.30	1.34	1.31	1.70	1.42	1.02	1.54	1.36
	Preprocessed α	**0.47**	0.62	0.36	0.59	**0.66**	**0.39**	**0.88**	**0.77**	0.38	**0.75**	**0.82**
	Online estimation of α	0.73	0.82	0.68	0.81	**0.85**	**0.72**	0.98	0.86	0.68	**0.89**	**0.93**
Lab error	BRDF	14.9	6.5	10.5	2.80	1.10	22.9	2.00	3.85	**0.36**	5.26	3.50
	Light	0.14	1.11	**0.25**	**0.78**	7.50	0.23	3.60	1.16	15.9	2.00	2.28
	Balance heuristic	0.14	0.81	0.31	0.91	0.77	0.28	1.04	1.21	0.47	1.65	1.60
	Power heuristic	**0.11**	0.83	0.31	0.93	0.78	0.28	1.06	1.22	0.46	1.70	1.59
	Maximum heuristic	0.12	0.95	0.31	0.99	0.85	0.28	1.30	1.37	0.45	1.98	1.71
	Combined estimator	0.13	**0.78**	0.29	0.86	**0.74**	**0.26**	**1.02**	**1.15**	0.44	**1.60**	**1.56**
Time (sec.)	BRDF	**159**	138	**40.6**	55.1	33.7	**16.9**	59.1	36.5	49.2	354	91.0
	Light	300	**103**	147	**41.8**	**18.6**	33.8	**41.7**	**34.2**	**27.1**	**178**	**53.0**
	Balance heuristic	238	124	95.5	48.9	28.0	25.5	52.0	36.2	40.0	281	74.0
	Power heuristic	240	122	95.7	48.8	27.1	24.9	49.9	35.9	39.2	275	74.0
	Maximum heuristic	239	122	95.4	48.1	26.0	24.5	49.5	34.9	38.5	264	71.1
	Combined estimator	251	128	107	45.8	28.3	26.9	51.1	34.5	40.9	281	72.0
Efficiency	BRDF	0.02	0.00	0.01	0.08	0.73	0.14	0.94	0.02	**10.3**	0.00	0.12
	Light	**1.51**	**2.24**	**12.4**	**2.24**	0.15	1.03	0.12	**1.26**	0.00	0.81	0.60
	Balance heuristic	1	1	1	1	1	1	1	1	1	1	1
	Power heuristic	0.89	1	0.99	0.90	0.94	0.89	0.93	0.94	1	0.93	0.93
	Maximum heuristic	0.81	0.94	1.00	0.87	0.85	0.89	0.68	0.78	1.02	0.73	0.82
	Combined estimator	1.43	1.18	1.34	1.45	**1.23**	**1.49**	**1.10**	1.29	1.44	**1.19**	**1.19**

Table 5.6.1: Relative average variance and Lab error, rendering time and relative efficiency for each scene and for different balancing methods. We use $N = 256$ samples per light source for each method. Combined estimator refers to the one introduced in Equation 5.16 where we use half of the total samples to evaluate α ($M = 128$) and the remaining half to evaluate the radiance with the estimated α ($N = 128$). The relative variance is computed as the variance of each strategy against the variance of the default balance heuristic. The Lab error is computed against the reference solution. The relative efficiency is computed as the efficiency of each strategy against the efficiency of the default balance heuristic. All timings were measured on an PC with an Intel CPU (i7-3930K@3.2 GHz).

al. [SHSK19] focus on optimizing sample allocation.

Sbert et al. demonstrate that using successive optimization of α converges toward the optimal strategy, still for direct lighting only. Their method optimizes the value of α through a numerical procedure using the Newton-Raphson method []:

$$\alpha_{k+1} = \alpha_k - \frac{V'_\alpha[F](\alpha_n)}{V''_{\alpha\alpha}[F](\alpha_k)} \quad \text{with } k \in [0..n_k] \text{ and } n_k \in \mathbb{N}^* \tag{5.17}$$

where $\mathrm{V}[F]$ is the variance of the estimator F, and where V'_α and V''_α are respectively the first and second derivatives according to parameter α. The initial value α_0 of this optimization procedure, is computed with Equation 5.12. Exact expressions for the derivatives can be found in the supplemental material of [16]. Sbert et al. also reduce further the interval size for α ([0.25, 0.75] vs. [0.025,0.975]), which prevents to use only a single strategy.

5.8 An improved version

In this section we present preliminary results for an improved version (published in [16]) of our per-pixel variance reduction algorithm. All results presented in this section are implemented (cf. Section 5.8.2) in our in-house rendering engine called Malia [27].

The main limitation of our method and the one introduced by Sbert et al. [SHSK19] are their inability to deal with scenes where indirect lighting becomes dominant.

A first and straightforward extension to global illumination is to reuse the α value estimated from the direct illumination for the first ray-bounce. For the other bounces, a random direction is selected with IS according to the BRDF and direct lighting is systematically estimated as in [Kaj86]. However, this is inefficient for most cases as illustrated in Figure 5.8.1. Indeed, this approach decreases the efficiency:

- slightly for open scenes, lit by an environment map, as most light paths correspond to direct lighting and only a small number of them consists in indirect paths.
- significantly for closed scene; the efficiency drops due to indirect light paths being widely under-sampled since light sampling as been identified as the best strategy.

However, note that with this approach, our original method ([15]) to compute α performs better than Sbert et al. (for $n > 1$) because α is more agressive clamped ([0.25, 0.75] vs. [0.1,0.9]) thus preventing the choice of a single strategy, which is not appropriate in the context of global illumination.

Another approach is to account for the radiance from *indirect* lighting contributions when computing α instead of the radiance issued from direct lighting only. This means that we still only use α to guide the samples allocation only for the first bounce of the light path, but this value will now include information regarding the full light path. Note that we cannot simply compute α per bounce per pixel as we would be mixing spatially incompatible information. By doing so, we can extend the previous methods to increase the efficiency in the context of global illumination (cf. Figure 5.8.2). As in the previous solution (cf. Section 5.2), this estimator converges toward a similar optimum independently of the initial choice of balancing.

Number of iterations vs. Number of learning samples

Obviously there is a tradeoff between the number of samples used for the learning phase and

the number of iterations during the optimization. Furthermore, this tradeoff depends on the scene. Nevertheless, as illustrated in Figure 5.8.3, the relative efficiency (compared to the balance heuristic) is always increased when learning the optimal α value with global illumination instead of direct lighting only.

5.8.1 Three Strategies Balancing

We now extend our improved method to include more than just light or BRDF sampling strategies. At the moment we only tested the simple case with three strategies at the same time. As a third strategy, we choose uniform sampling on the hemisphere: it ensures that the third strategy is not correlated[1] to the other two while ensuring the exploration of the full space of light paths. With three strategies (light, BRDF and uniform sampling), the pdf relative to defensive sampling becomes:

$$\mathrm{p}_{\alpha,\beta} = \alpha \mathrm{p}_{\mathrm{BRDF}} + \beta \mathrm{p}_{\mathrm{light}} + (1-\alpha-\beta)\mathrm{p}_{\mathrm{uniform}} \tag{5.18}$$

with $\alpha, \beta \in [0,1]; \alpha+\beta \leqslant 1$. In this case, minimizing the variance means that we are looking for α, β such that:

$$\nabla \mathrm{V}[F](\alpha,\beta) = \begin{bmatrix} \mathrm{V}'_\alpha[F](\alpha,\beta) \\ \mathrm{V}'_\beta[F](\alpha,\beta) \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \end{bmatrix} \tag{5.19}$$

$$\det(\mathscr{H}(\mathrm{V}[F](\alpha,\beta))) > 0 \tag{5.20}$$

$$\mathrm{trace}(\mathscr{H}(\mathrm{V}[F](\alpha,\beta))) > 0 \tag{5.21}$$

where $\mathscr{H}(\mathrm{V}[F](\alpha,\beta))$ is the 2×2 Hessian matrix of the variance. The full formulation of $\mathscr{H}(\mathrm{V}[F](\alpha,\beta))$ is provided in Section 4.2 of the supplemental material of [16]. in which we demonstrate that assumptions of Equations 5.20 and 5.21 are valid in our case where the strategies are not correlated. The supplemental material also presents the general formulation with an arbitrary number of balancing strategies.

Therefore we only need to solve Equation 5.19, using the Newton-Raphson method:

$$\begin{bmatrix} \alpha_{n+1} \\ \beta_{n+1} \end{bmatrix} = \begin{bmatrix} \alpha_n \\ \beta_n \end{bmatrix} - \mathscr{H}^{-1}(V[F](\alpha_n,\beta_n))\nabla V[F](\alpha_n,\beta_n) \tag{5.22}$$

As in [15] and [SHSK19], all these derivatives are estimated by Monte-Carlo integration using the same rays as the ones used to estimate the radiance at the same time.

Moreover, as we did for two strategies, we must clamp $\alpha+\beta$ in the range $[0.1, 0.9]$ to ensure that we keep exploring all light paths.

Note that our third strategy (i.e., uniform sampling) is not expected to perform well and we expect that in most cases, $\alpha+\beta \geqslant 0.9$, yielding to an allocation close to the one we obtained

[1] If the third strategy is correlated to the others, a bias would be introduced.

with only two strategies. However, our main goal is to highlight that our technique is not limited to two strategies and can identify the local optimal strategy between more than two, as illustrated on Figure 5.8.4. We can see that, even when adding a simple strategy, our approach still increases the efficiency since it gives higher weights to near-tangent contributing paths. More precisely, in Figure 5.8.4c, $\epsilon_{rel,2strat} = 124\%$ represents the relative efficiency of our method with 3 strategies relatively to using the balance heuristic with 2 strategies (each strategy with the same number of samples), whereas $\epsilon_{rel,3strat} = 40\%$ represents the relative efficiency of our method with 3 strategies relatively to using balance heuristic with 3 strategies. In another words, our 3-strategies method performs better than the balance heuristic with 2 strategies!

5.8.2 Malia Implementation Overview

Malia is a GPU-based path-tracing renderer based on Optix that supports spectral rendering and multiple importance sampling. Compared to a classical OpenGL/DirectX engine, Optix supports visibility query on the GPU, which drastically improves performance compared to offline (CPU) engine (e.g., PBRT).

5.9 Short-Term Future Work

As a first experiment, we have limited the balancing to the first bounce. The resulting gains lead us to believe that it can be generalized to any bounce. For this purpose, a spatial structure (e.g., with a caching approach [GKPS12]) must be developed and trained. With such an extension, we may also learn how much we can limit the number of shadow rays for direct lighting at any bounce. Most approaches that aim at efficiently resolving complex light paths (e.g., [BJ19]) could benefit from our solution, in order to better identify these complex paths.

Using a simple uniform sampling as a third strategy is not the best choice. We have shown in our preliminary tests some gain in efficiency. This demonstrates that considering more tailored strategies, forcing to explore the space not explored by the classical ones, may result in more significant gains. For example, sampling BRDF in the retro-reflection direction or in the tangent one, will conserve the orthogonality of the third strategies and should improve the efficiency. We could even add these two sampling strategies to obtain a 5-strategy balancing scheme. Moreover, using cache-based pdf such as in [HEV+16b],may also result in more significant gains.

Another possible strategy that could be included, is one that decides to evaluate or not direct lighting when propagating a path. For scenes with very complex geometry, where most likely shadow rays will be blocked, avoiding the visibility query, could improve the overall efficiency of the method.

Our generalized approach to more than two strategies, combined with the fact that the convergence is independent from the initial balancing strategy, opens a new field of studies to learn the best strategy in the context of Multiple Importance Sampling.



Scene 3
Diffuse BRDFs
Three area light sources

Scene 4
Diffuse BRDFs
High-frequency environment map

Scene 5
Highly specular BRDF
High-frequency environment map
+ 1 point light source



Scene 6
Diffuse BRDFs
Area light source
+ 1 point light source

Scene 7
Glossy BRDF
Low-frequency environment map

Scene 8
Mirror BRDF
High-frequency environment map



Scene 9
Glossy BRDFs and complex visibility
Low-frequency environment map

Scene 10
Diffuse BRDFs and complex visibility
Low-frequency environment map
+ 1 point and 1 directional light sources

Scene 11
Measured glossy BRDFs
High-frequency environment map

Figure 5.6.2: Different test scenes, ranging from diffuse (scenes 3, 4 and 6) to specular (scenes 8 and 2) materials including complex visibility (scene 10) with different types of illumination such as area light sources (scenes 1, 3 and 6) and environment maps. Since only one sample is required for point and directional light sources, they are added in a post-process and removed from the variance computation.

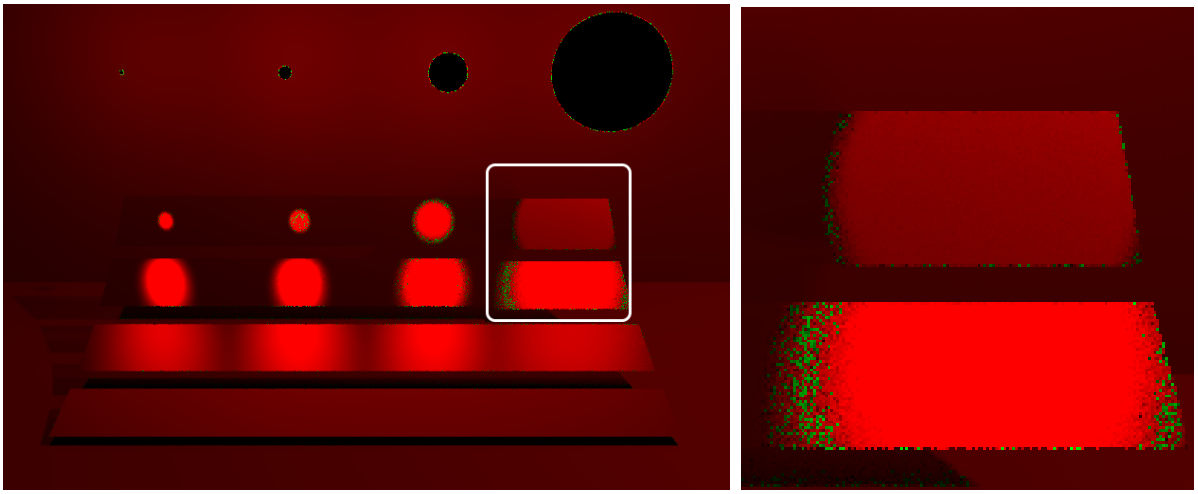


Figure 5.6.3: Visualization of per-pixel variance. These images are computed as the difference of variance between the default balance heuristic and our method. **Red** corresponds to pixels where our method is better whereas **Green** corresponds to those where the balance heuristic is better. **Black** means that none of the methods outperforms the other (i.e., variances are equal).

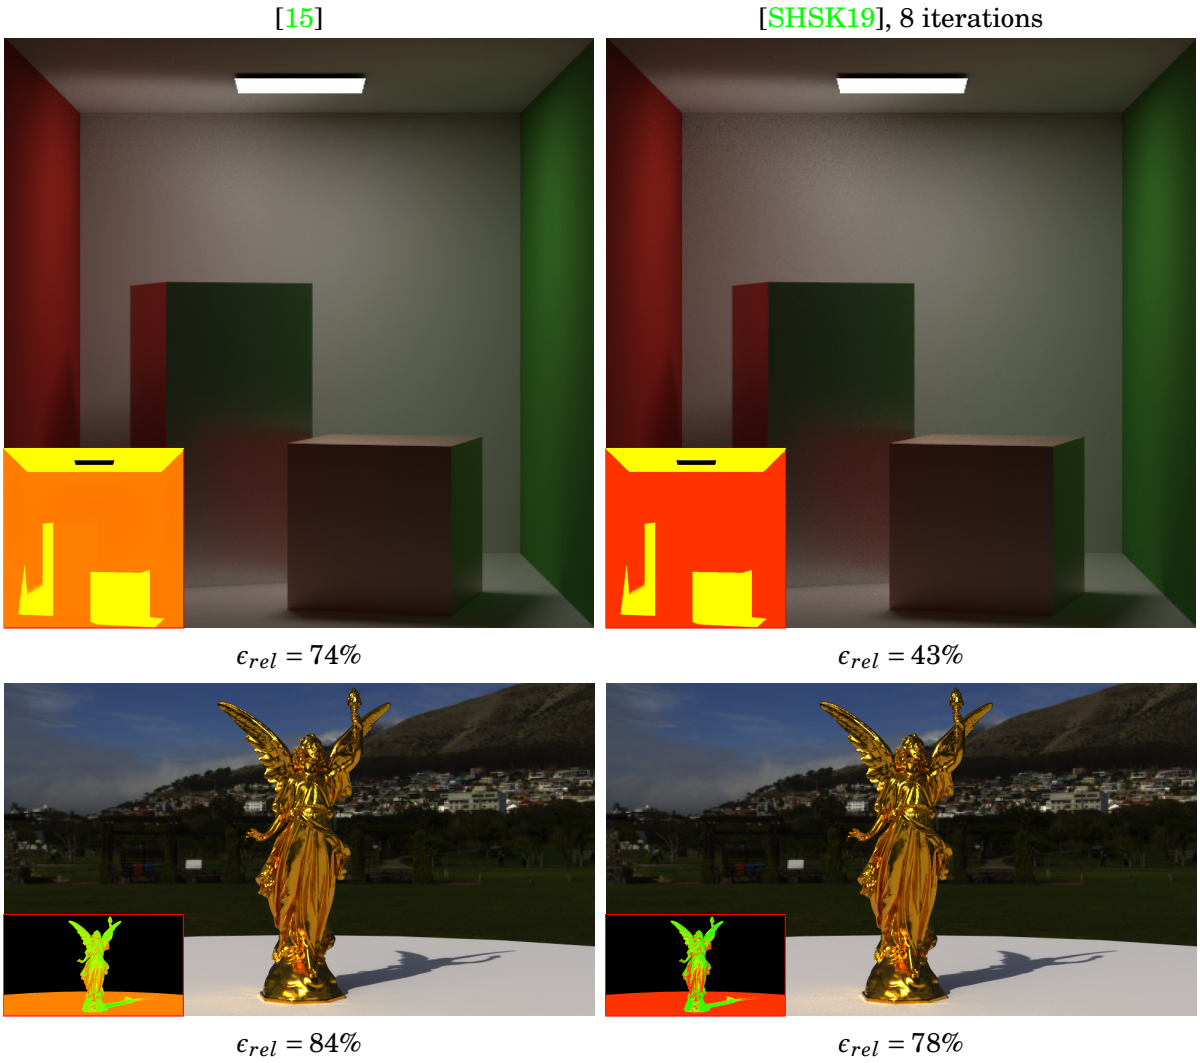


Figure 5.8.1: Renderings with Global illumination using precomputed α-maps, obtained with [15] and [SHSK19] only using direct lighting and 256 learning samples per iteration. **Top Part:** Closed scene with direct and indirect illumination. The relative efficiency drops as the variance increases significantly due to indirect paths being under-sampled. **Bottom Part:** Open scene dominated by direct illumination. The efficiency drop is less important as most light paths are indeed direct ones, but still present because inter-reflections are poorly sampled.

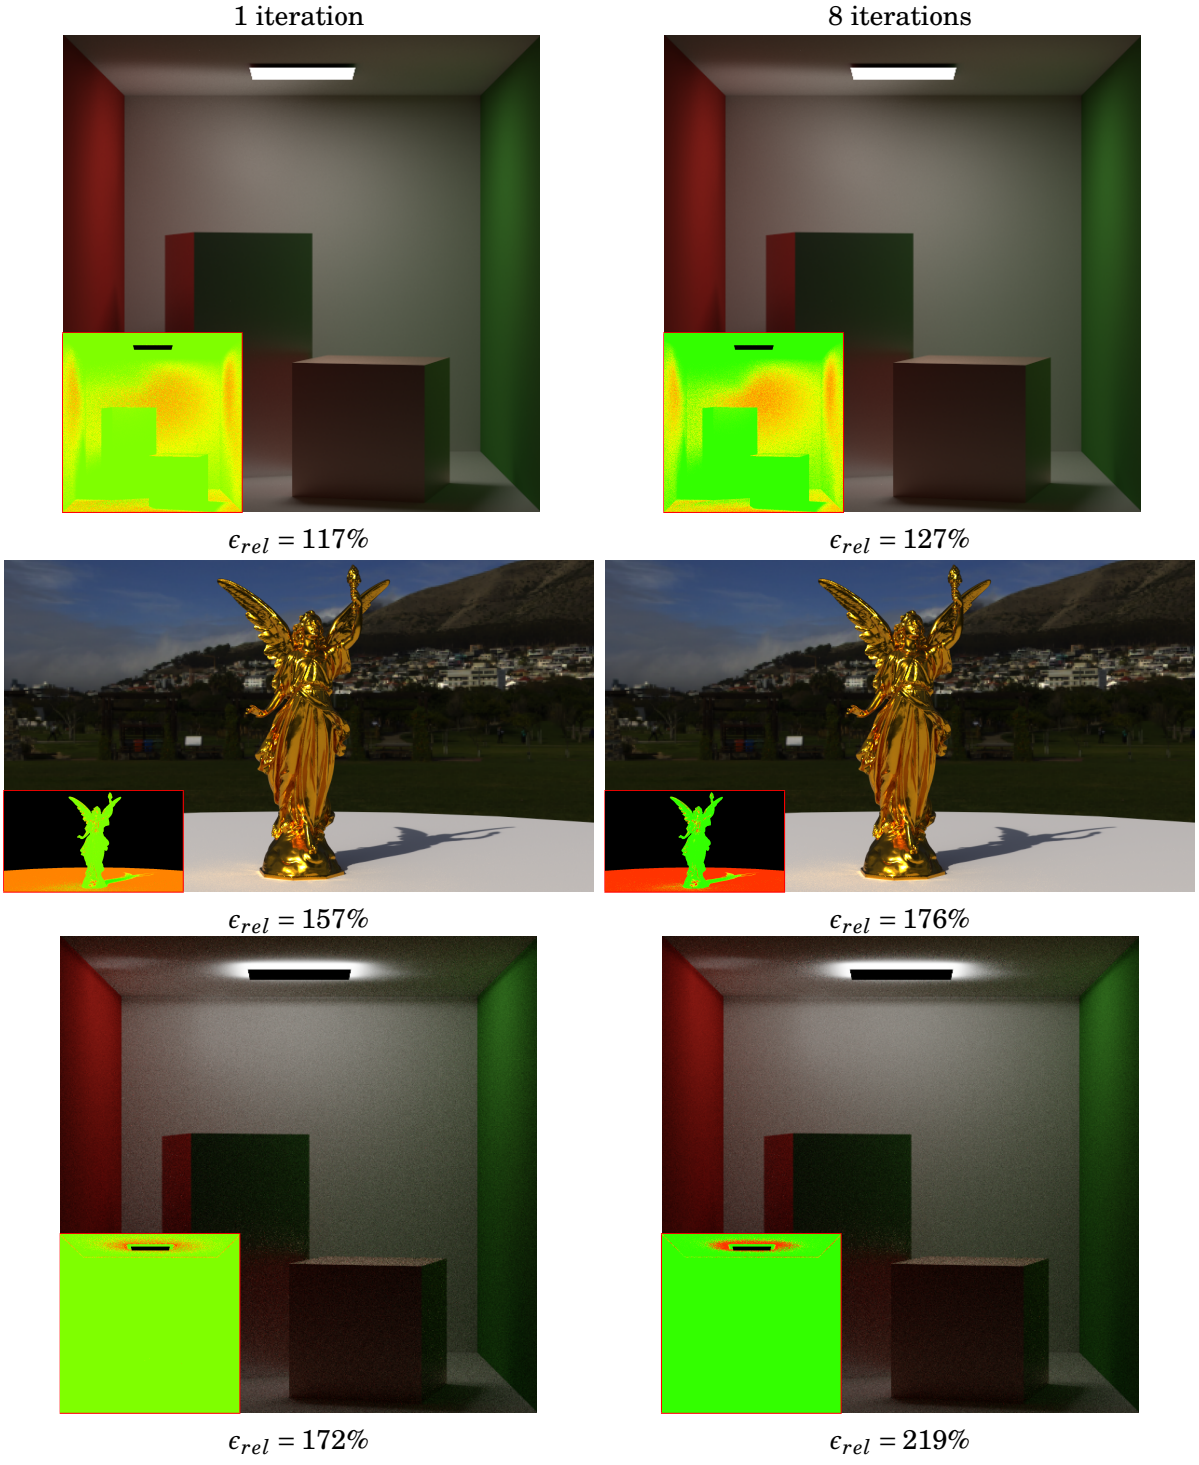
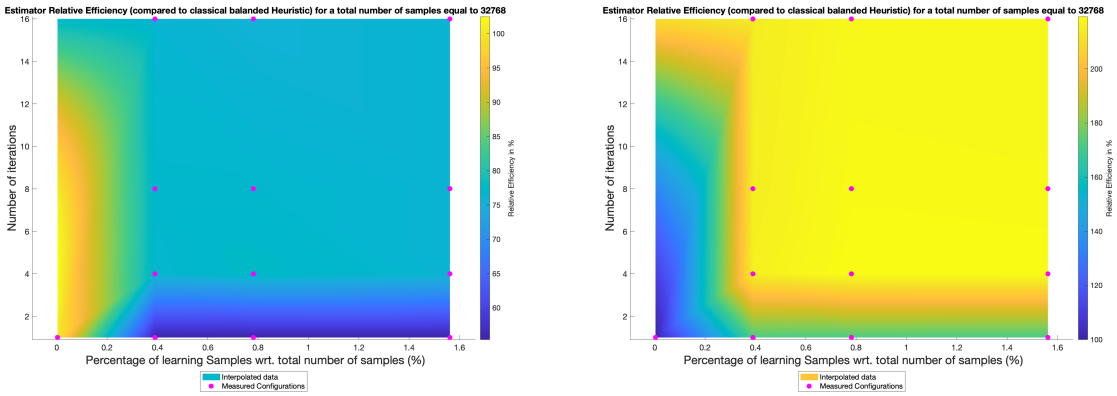
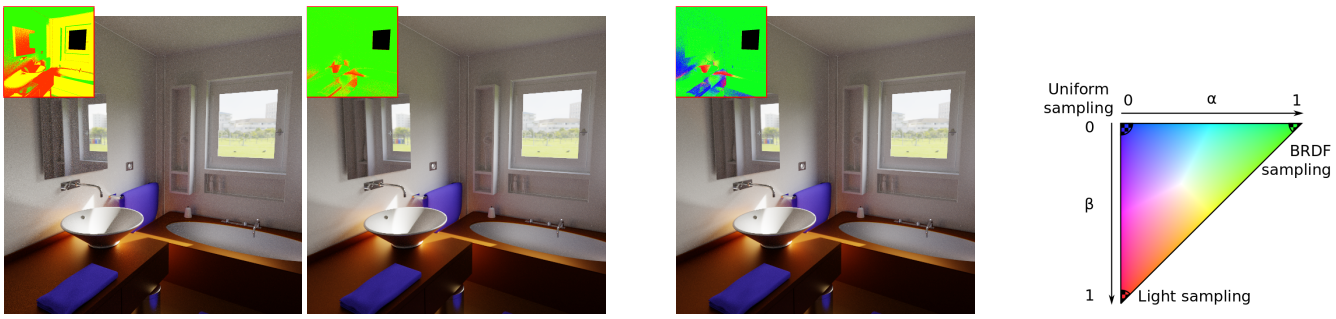


Figure 5.8.2: Comparison between one iteration and 8, using global illumination to compute α. Rendered with 256 learning samples per iteration. **Top row:** In areas where indirect contributions are greater that direct ones, BRDF sampling is identified as optimal whereas light sampling is optimal where direct contributions are greater. This results in a decreased variance at nearly equal rendering time. **Middle row:** The behavior is close to the one in Figure 5.8.1 except in areas where there are many inter-reflections (e.g., concavities), and thus BRDF sampling is optimal, leading to a reduced variance and increased efficiency. **Bottom row:** BRDF sampling is nearly a global optimal strategy, yielding a significant gain of efficiency.



(a) (b)

Figure 5.8.3: Modified Cornell box with strong indirect lighting (cf. last row of Figure 5.8.2). Relative efficiency of the Monte-Carlo estimator compared to the balanced heuristic with equal allocation between light and BRDF samples. The X-axis is given per iteration. **(a)** α is computed with direct illumination only vs. **(b)** global (direct+indirect) illumination. **Using our improved version of learning (b) can achieve 200% efficiency against the balanced heuristic whereas the previous techniques ([15, SHSK19]) achieve only 90%.** These timings were measured using an Nvidia® RTX 2080 graphics card.



(a) $\epsilon_{rel} = 62\%$ (b) $\epsilon_{rel} = 144\%$ (c) $\epsilon_{rel,2strat} = 124\%$; $\epsilon_{rel,3strat} = 140\%$ (d)

Figure 5.8.4: Optimizing sample allocation comparisons for the first bounce between (a) two sampling strategies (light and BRDF) using [SHSK19], (b) two sampling strategies using **our solution**, (c) three strategies (by adding a uniform sampling strategy). (d) Color map used to display the sample allocation (top left corner of each image) for each strategy. Images are rendered using 512 learning samples and 8 iterations to compute α. The efficiency is given relative to the reference computed with the balanced heuristic computed and the same number of samples for each strategy.

Part IV

Future Directions

CHAPTER 6

FUTURE DIRECTIONS

6.1 Main Objectives

As stated in the introduction, computer-generated images are now experiencing strong growth in both manufacturing industry and human sciences. These new use cases raise new challenges and new scientific issues at the interface of instrumental optics and digital technologies. In parallel, humanities and social sciences (SHS) are facing new societal challenges with the preservation of heritage, but also new opportunities thanks to the progress of opto-numerical technologies and 3D in general. Digitizing ancient cultural heritage objects helps their preservation. 3D digital tools allow us to better understand how our ancestors perceived objects or to determine the different stages of degradation that led them to their current appearance.

Since the beginning of the industrial era, prototyping has always been an important preliminary s tep before mass production. With the emergence of computer graphics and advances in computing power, the industry is progressively moving from tangible prototypes (e.g., mock-ups made with the same manufacturing processes as the final object) to fully digital prototypes (virtual mock-up) with the stated objective

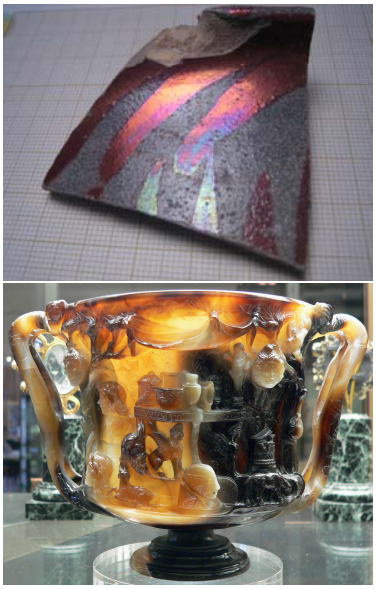


Figure 6.1.1: Examples of ancient materials with a complex appearance. **Top:** Lustrous ceramic from the 9th-10th century (AD) from Mesopotamia. The angularly varying colour effects can be explained by the presence of nanoparticles. Photograph from Reillon [Rei08]. **Down:** Ptolemaic cup (First century BC) with a translucent material that diffuses light into the volume. Photograph from Cabinet des Médailles.

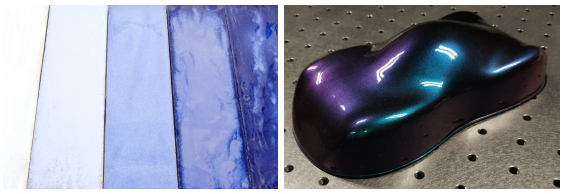


Figure 6.1.2: Examples of industrial materials with a complex appearance. **Left**: glazed ceramic. **Right:** paint with colored effect

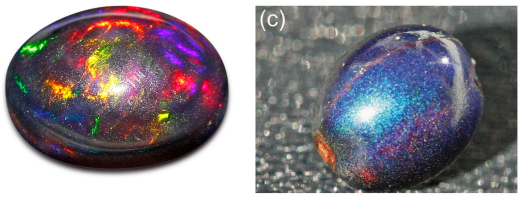


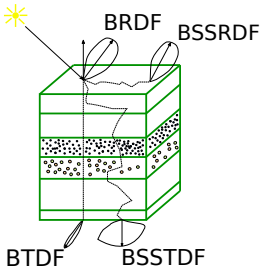
Figure 6.1.3: Examples of natural materials with a complex appearance. **Left**: opalescent stone. **Right**: colors obtained by nanostructuring of cellulose

of reducing production costs. Over the last ten years, the growth of 3D printing technologies allowed the development of (low-cost) virtual prototyping methods that make it possible to know upstream how far the realization will be made possible; in particular for the object shape. Beyond the shape, predicting the appearance of a virtual prototype remains a challenge with a strong applicative potential for different industrial sectors (e.g., furniture, textile, construction, architecture...)

Since appearance is the result of the interaction of three elements: the shape (or volume), its material and the lighting environment. This combination poses many challenges when it comes to measuring, simulating, modelling or predicting appearance. This is due to the fact that the main physical quantities related to appearance are high-dimensional, ranging from 3D, (e.g., isotropic BRDF) to 12D (e.g., BSSRDF). This is all the more true as ancient objects (cf. Figure 6.1.1) and industrial creations (cf. Figure 6.1.2) or natural (cf. Figure 6.1.3) objects, already present a great wealth of appearances, and tomorrow's creations (e.g., structured by nano-technologies), whose appearance can also be inspired by nature or ancient techniques, will enlarge the space of possible appearances.

Therefore, my long term objectives can be summarized as follows:

Acronym	Dimension	Signification
BRDF	4	Bidirectional **Reflectance** Distribution Function
BTDF	4	Bidirectional **Transmittance** Distribution Function
BSSRDF	8	Bidirectional **Subsurface Scattering Reflectance** Distribution Function
BSSTDF	8	Bidirectional **Subsurface Scattering Transmittance** Distribution Function



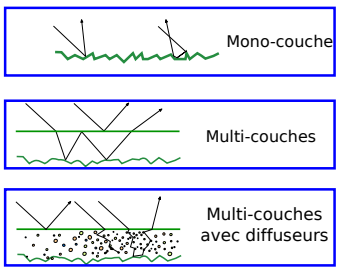


Figure 6.1.4: **Left**: Nomenclature of the main radiometric quantities for a homogeneous material on the surface. **Middle**: Example of a multilayer material showing the different radiometric quantities involved. **Right**: Three levels of complexity for a multilayer material. Although the radiometric quantities are already high dimensional, in general, spatial, spectral, temporal and polarisation dependencies should be added to take into account all optical phenomena.

1. To establish the links between optical properties at different scales and the measurement of appearance by digital sensors.
2. To develop new smart, efficient and metrologically validated measurement devices.
3. To predict the appearance (i.e., *inverse design*) of manufacturable objects and virtually restoring the appearance of ancient objects.

In order to achieve these objectives, my research project is organized around three axes:

- Axis 1: Appearance measurement with quantified uncertainties.
- Axis 2: Light transport simulation and appearance synthesis.
- Axis 3: Inverse models for appearance prediction.

By developing qualified measuring instruments, Axis 1 serves both to acquire the appearance for preservation purposes and to validate the various models developed in Axis 3. Axis 2 allows for the analysis and understanding of physical phenomena at the wavelength scale of light and the study of their links with the macroscopic scale in order to develop the predictive models of Axis 3.

6.2 Axis 1: Appearance measurement with quantified uncertainties

6.2.1 To validate the predictions of models

Whether they are used to capture the appearance of the real world or to validate the appearance of predictive models, the measurement devices will need to:

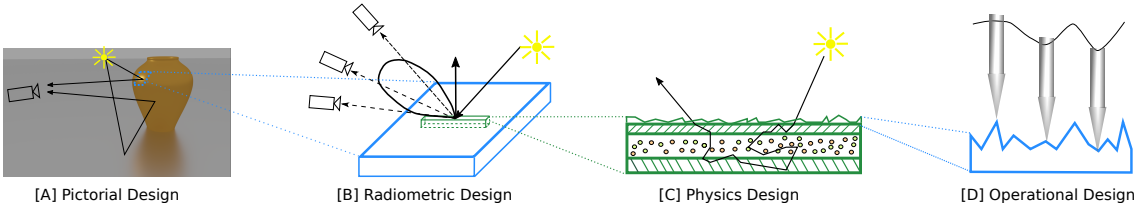


Figure 6.1.5: Overview of the proposed methodology for the prediction of the appearance of multi-layered materials in order to perform inverse design. To control the final appearance of an object, one choose to go back to the physical parameters of the layers. The transition from [C] to [B] to [A] can be done by measurement (Axis 1) or by numerical simulation (Axis 2). The transition from [A] to [B] to [C] represents the inverse design and requires models (Axis 3). The transition from [C] to [D] represents the inverse design operation of manufacturable objects and depends on the manufacturing process. The VIDA ANR Project (2017-2021) focuses on the transition between [B] and [C]. In the near future, I will focus on the transition from [B] to [C] to [D] thanks to the help of partners specialized in material fabrication.

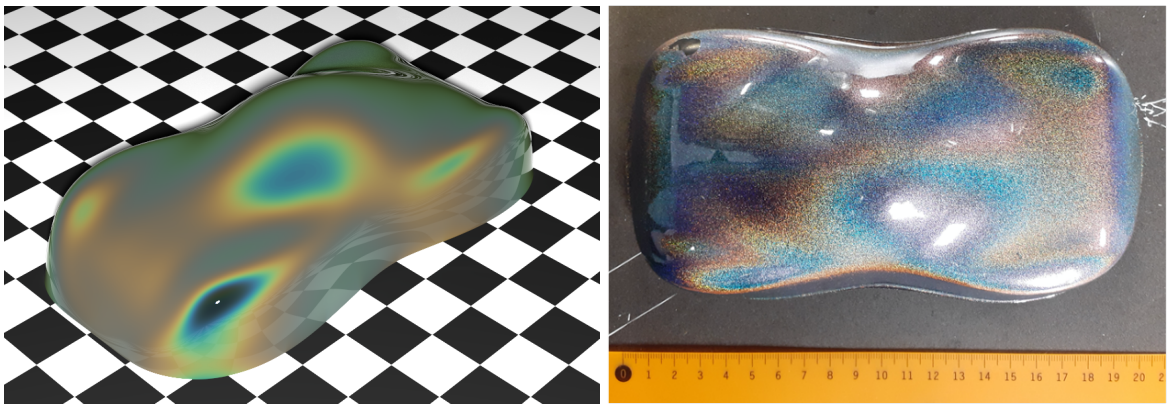


Figure 6.2.1: **Left:**. Light Transport simulation of a speedshape coated with correlated nanoparticles on a substrate (see Figure 1.4 left for more details). **Right:** A real speedshape coated with iridescent painting distributed by Stardust company.

(i) be capable to perform dense measurements,

(ii) being qualified (e.g., uncertainties) and calibrated against level 1 instruments developed in national metrology centres (e.g., LNE gonio-spectro-reflectometer [LBGO15]).

Obtaining dense measurements will make it possible to avoid making strong assumptions about the measured phenomena (e.g., spatial, directional or frequency homogeneity) while the measurement uncertainties will allow a relevant evaluation of the developed models. As we have seen in Chapter 2, the high dimensionality (see Figure 6.1.4) of the various radiometric quantities poses a challenge in terms of managing the volume of data generated, the complexity of hardware implementation, and the acquisition time. Very often, the most accurate systems are extremely slow (several days of measurements) for a very small sample size. In my opinion, it is therefore appropriate to proceed in a hierarchical manner for the development of appearance measurement systems.

Therefore, in the short term, I intend to develop a simplified, yet accurate and calibrated devices to measure appearance directly in image (or sensor) space (see [A] in Figure 6.1.5). For example, by fixing the illumination, and limiting the number of viewpoints, it will be possible to quickly validate predictive models of the appearance of a manufactured object. Validation in image space does not guarantee that the predicted quantities are accurate in radiometric space. Thus, in the long term, my project also aims to continue the development of dense and fast measurement instruments. This will be done in particular through optical and numerical co-design.

The next application of qualified measurements will take place in a collaboration with chemical colleagues at ICMCB. Instead of using lithography, they will chemically create correlated structures deposited on a substrate in order to obtain the same type of appearance as those we have published. By comparing the BRDF obtained with this new fabrication process to previous

measurements we will be able to validate it at two scales. By measuring the BRDF when the material is deposited on a small (1cm) planar surface but also on a large curved surface (20 - 30 cm span). The interest of this chemical process is to be able to carry out, at a lower energy cost, a coating on a large surface. Finally, it will be interesting to compare the measurements with more conventional coatings (diffractive flakes contained in the so-called iridescent paints), which are currently used in industry (see Figure 6.2.1). The goal would be to achieve the appearance of existing paints as well as new effects such as the reduction of specular reflection that we obtain with *correlated* nano-particles.

6.2.2 For Cultural Heritage Preservation

Although some ancient objects can be moved to the laboratory, this is not always possible for reasons of space, security or fragility. It is therefore necessary to develop and perfect acquisition devices capable of performing in-situ measurements. Although some optical measurement devices are transportable, devices with the capacity to perform on-site appearance measurements are often limited. Some rely on strong assumptions (e.g., [RWS+11]); others remain limited in terms of angular or spatial resolution (e.g., [ACGO18, AWL15]). Finally, portable devices developed for computer graphics applications share the same limitations in terms of metrological qualification as those developed in the laboratory.

In the long term, I would like to develop portable instruments capable not only of being fast in terms of measurements but also in terms of feedback to the experimenter. Indeed, an on-site acquisition can rarely last more than a day and it is important to guide the experimenter during the acquisition. This can be achieved by using assumptions about the material of the object to be scanned but, unlike previous approaches, these assumptions will be determined by a dense measurement of a sample taken in-situ beforehand (where possible). Without a priori, other interpolating or extrapolating approaches (e.g., [DWT+10]) are to be developed. These methods re-combine sparse measurements to generate a dense measurement. By injecting insights from the field and from experts in humanities and social sciences or material sciences, I plan to improve the current techniques (e.g., [DWT+10, ZCD+16]).

6.3 Axis 2: Light transport simulation and appearance synthesis

The aim of Axis 2 is to enable a better analysis and understanding of light-matter interaction phenomena in order to predict the radiometric responses and appearance of objects. This will allow to establish links between the physical parameters (cf. [C] in Figure 6.1.5) and the macroscopic appearance (cf. [A] in Figure 6.1.5) in order to develop direct and inverse models (cf. Axis 3). To this end, my objective is to develop predictive simulation tools at the wavelength scale to calculate radiometric responses and, at the macroscopic scale, to predict the appearance of an object.

6.3.1 Radiometric simulation

Given the complexity of materials both in their composition (volume, layer stacking) and in their surface structuring, it makes sense to develop specific simulation tools whose results are then combined to obtain the final radiometric response of the material. In the short term, my objective is to develop a simulator for multilayer materials that takes spectral and polarisation dependencies into account, while achieving high performance in terms of calculation time. A first CPU-version of a spectral simulator [28] has already been implemented in the context of the Vida project, and an optimized one will soon be completed thanks to our recent work on efficient spectral rendering on the GPU [29]. In addition, the results of this simulator will be validated against simple cases for which analytical solutions exist but also against radiometric measurements of known samples.

In the medium and long terms, I intend to develop numerical tools that can be used beyond the validity assumptions (coherent regime) of the Radiative Transfer Equation (RTE) [Cha60]. My recent work with the nano-photonists of LP2N on structures with nano-particles deposited on a planar substrate cannot be solved via the RTE because of strong coherence effects. Despite the existence of a multitude of numerical methods for solving the electromagnetic equations (e.g., FDTD [Yee66], finite elements [ZTZ05], T-Matrix [MTM96]), they remain little used for large systems due to their computational time and memory cost. Although these computational methods do not explain the relationship between nano-structures and the final appearance (see Axis 3), it seems important to provide solutions in terms of optimisation (algorithms and models) and parallelism in order to obtain (predictive) reference solutions for this type of systems. All the more so as the alternative solution of manufacturing and measuring such samples remains, for the moment, difficult to achieve. Indeed, the manufacturing process (e.g., electron beam lithography) is time and energy consuming and the size of the fabricated sample remains very small. Consequently, measuring the BSDF of such small elements requires very specific instruments.

6.3.2 Predictive Rendering

Generating synthetic images is a field that has been in constant evolution for almost 50 years (cf. [PJH16]). In my research project (cf. [B] to [A] of Figure 6.1.5), I wish to bridge the gap that remains between images that appear realistic to us and predictive images on which uncertainty is established. Indeed, most of the developed renderers are not predictive because they do not take into account: (i) the spectral dependence of quantities (e.g., refractive index), (ii) fluorescence (change of wavelength) and polarization effects. Of course, taking into account these different effects, reinforces the computation time and memory space problem. Therefore, in order to obtain images (or animations) in a reasonable time, it is necessary to further develop solutions that reduce memory costs and computation times. The development of models (analytic or tabulated) is one of the possible solutions to reduce memory costs (cf. Axis 3) but approaches such as signal

compression (e.g., function bases [ZERB05], multi-dimensional decomposition [BOK11]) are also to be improved, in particular by pursuing the development of new parametrizations such as the one we have proposed (cf. [23]).

In addition, I intend to pursue the development of algorithms suitable for parallel architectures on graphics cards which are in constant evolution. In the continuity of existing work (e.g., [Vea98a, WND+14]) these new parallel methods will have to be coupled with new importance sampling methods. Importance sampling methods will allow to gain an additional order of magnitude in terms of efficiency and are more than ever necessary to speedup the convergence of the light transport simulation when taking into account supplementary physical phenomena such as spectral, fluorescence, polarization or light coherency effects.

6.4 Axis 3: Inverse models for appearance prediction

The objective of this research area is to develop new inverse models to:

(i) determine the optical properties of manufacturable objects from a desired appearance;

(ii) reconstruct the processes that degrade materials and alter their appearance in order to estimate their original constitution.

Solving the inverse problem from image space [A] to manufacturing operational space [D] is a difficult and ill-posed problem (see Figure 6.1.5). Thus, the methodology I propose is to develop a framework for direct and inverse appearance design without constraining to a particular manufacturing process. As the manufacturing processes are numerous and diverse, it is important, in a first phase, to model the links between physical quantities and the final appearance of the macroscopic object. In a second phase, the developed inverse models will be coupled and instantiated for specific manufacturing processes before being validated by measurement instruments (see Axis 1). In the future, the direct and inverse models that will be developed will be associated with effective representations such as simplified or application-constrained models (e.g., refractive index > 1) that provide feedback to the user in interactive time. Once an initial solution is found, a numerical optimisation process will instantiate a predictive solution (see Axis 2). This will enable design tools whose output will be related to the physical parameters of the materials and, unlike some previous work (e.g., [BAEDR08]), they will not be constrained to a fixed view or rely heavily on long pre-computations (e.g., [WDR11])

In the short term, I will continue to investigate multi-layered materials, which are very promising in terms of appearance, by increasing the complexity of their content (particle size and density) and their interfaces (smooth vs. rough) or even their periodicity such as 1D to 3D photonic crystals (cf. Figure 6.4.1), or semi-periodic structuring. Multi-layered materials are of interest not only for industrial applications but also for historical (e.g., [Rei08]).

As a first step, I will consider a single interface (ranging from perfectly smooth to very rough) with eventually a single substrate (see Figure 6.1.4 Top-right). Although the link between surface

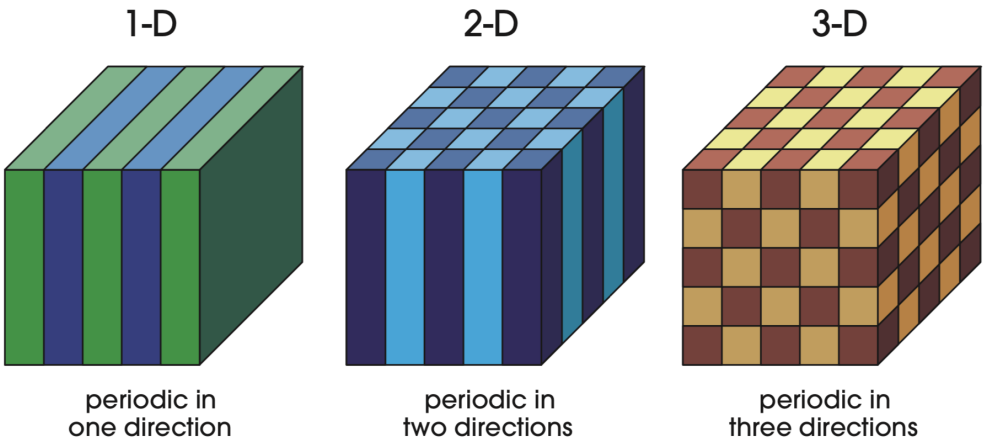


Figure 6.4.1: Photonic Crystals are periodical multi-layered systems where the index of refraction varies periodically in 1D or 2D or even 3D. Images from *Photonic Crystals, Molding the Flow of Light* [DGND08]. Observe that layered materials presented in Figure 6.1.4 are 1D and none-periodic systems.

properties (e.g., roughness) and optical response has been extensively studied (e.g., [Ric51b, KoCF06]), further evaluation of recent BRDF models (e.g., [22]) is still re quired in order to, ideally, establish a bijection with, eventually, the manufacturing parameters of machining devices. This evaluation can only be done by accurately manufacturing control materials and then measuring them just as accurately. With the help of industrial partners (e.g., Saint-Gobain, Valeo) to manufacture the samples and colleagues from the electromagnetic community (e.g., Inria Project Atlantis), we will be able to establish up to what extent BRDF models based on diffraction are sufficiently accurate.

We have already established a first inverse method to design the appearance for a stack of thick layers (see, Bati et al. [19]). Thanks to our analytical results, our method is very efficient in terms of computation. Beyond thick-film systems, I also intend to look at more complex systems that generate interferences.

Pertubated Thin-films without nano-particles

Such systems have industrial applications (cf. Figure 6.4.5 left) where they are considered as quasi-perfect when manufactured but they can also be found in biological realizations (cf. Figure 6.4.5 right) where they cannot be considered as perfect and where they can also be mixed with thick layers. Therefore, it is also necessary to be interested in the optical stability of such systems in order to establish which physical characteristics are sensitive (resp. invariant) and modify strongly (resp. weakly) the final optical response. Of course, this can be done by disrupting the different physical characteristics and comparing the variance of the different optical responses, obtained through numerical simulation.

To close the loop with a manufacturing process, members of the Manao team and I have started to collaborate with Mélanie Bay (former PhD Student at Cambridge University), who studied

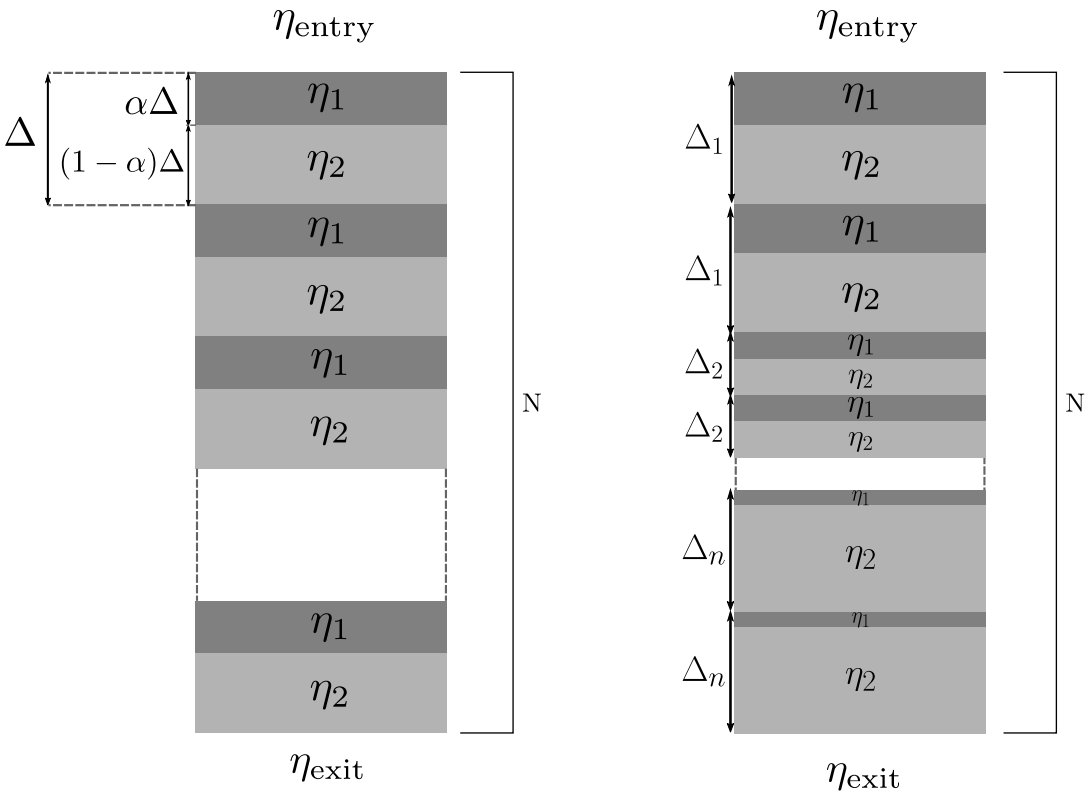


Figure 6.4.2: The different types of 1D periodic systems. **(Left)** A Bragg Mirror system is structure which consists of alternating two optical materials with different indexes of refraction η_1 and eta_2 and thicknesses ($\alpha\Delta$ and $(1-\alpha)\Delta$ with $\alpha \in [0,1]$). **Right)** A clustered system is the concatenation of multiple Bragg mirrors with different thicknesses. When the thickness of a chirped system varies continuously, the system is said to be "clustered".

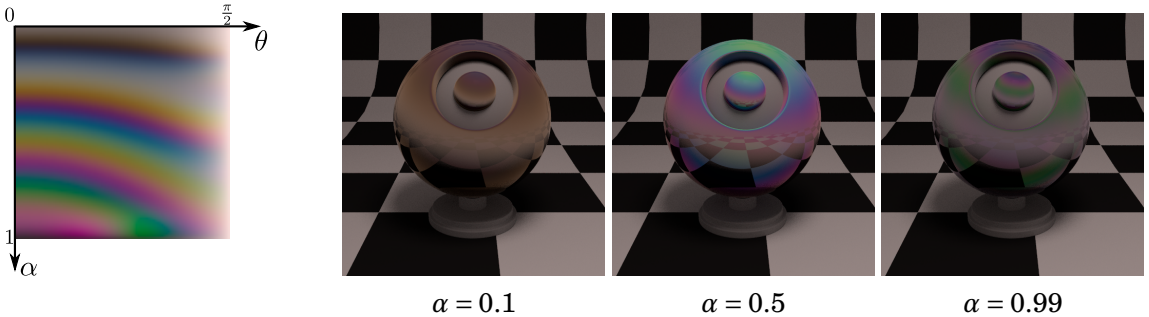


Figure 6.4.3: Predicted appearance of a Bragg stack which physical properties are: $\Delta = 700$nm, $N = 5$ and indexes of refraction are $n_1 = 1.4$ and $n_2 = 1.7 + 1i$. **Left:** Visualization of the different colour that can be achieved for each elevation angle (abscissa) and for each α value (ordinate). **Right:** Rendered spectral images for different values of α. Thanks to Pierre Hubert-Briere for these images.

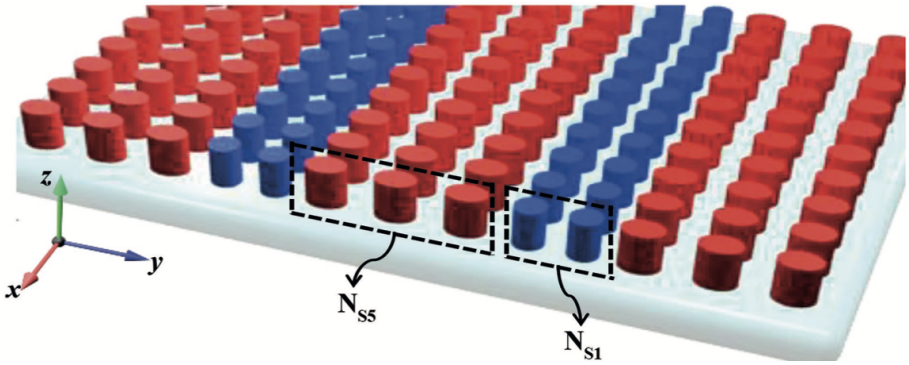


Figure 6.4.4: Periodic systems with elliptical nano-particles. Image from [YLL+18].

(cf. [Bay21]) and fabricated multilayered-materials that are obtained through the structuring of hydroxypropyl cellulose (a bio-sourced and water-compatible polymer). For this specific type of systems, we aim at developing an efficient macroscopic BRDF and inverse it to establish a complete inverse design method for bio-inspired materials.

Thin-films with nano-particles of arbitrary shapes

This type of system consists of a substrate on which a layer of nano-particles is deposited. The geometry of the nano-particles can be of any shape and the particles can be deposited in a periodic or semi-periodic way. An example of this type of system is presented in Figure 6.4.4. This type of structures has been seldomly studied in terms of BSDF appearance (only in terms of average color) despite the fact that the are promising in terms of expressiveness.

Volumic Systems with 3D Correlation

In line with our work [30] and the work of Guo et al. [GJZ21] and in collaboration with K. Vynck, we intend to study the appearance of 3D correlated systems. This require to simulate the light transport inside a volume according to the RTE but with new phase function models and scattering and absorption properties. Contrary to the work of Guo et al. [GJZ21], we intend to study systems that are *not clustered* and that can even exhibit disordered hyperuniformity (e.g., [HS13]).

Finally, no matter the complexity of the system, my approach remains the same and requires:

(i) To solve the electromagnetic (EM) problem through simulation either in partnership with the Inria ATLANTIS team, or by using existing (EM) solvers such as Smuthi [ECT+21] or RCWA [MG81] (e.g., S^4 [LF12]) according to the studied physical system. Typically, RCWA is meant for periodic or semi-periodic systems whereas Smuthi is for non-periodic systems.

(ii) To use the resulting tabulated solution from the EM solver to render a ground truth reference BSDF as well as a rendered image.

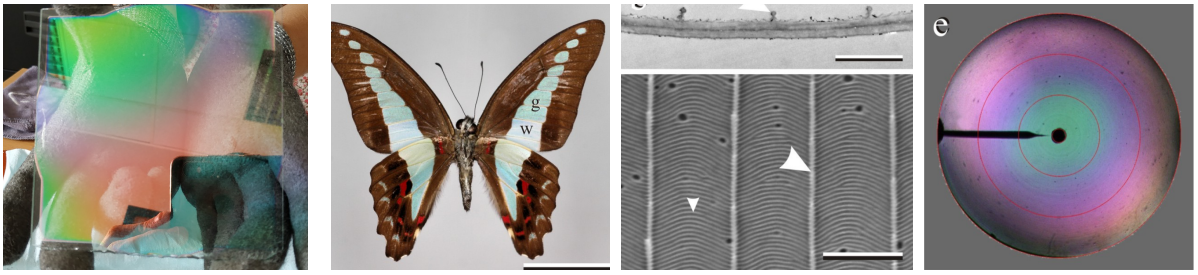


Figure 6.4.5: **Left:** Manufactured glasses from Saint Gobain with dozen of thin layers. Each layer thickness is about 10 nanometers and their combination produces colored interferences. Their realization requires a very fine control of the manufacturing process. **Center Left:** Natural example of thin layers that are not perfectly aligned and present irregularities: Graphium Sarpedon butterfly (images from [Sta14]). On the wings of the butterfly, unpigmented glass**(g)** or white **(w)** scales are present. **Center Right:** Electron microscopy scan of the part (g) revealing a complex and imperfect thin film structure. **Right:** Scattering measurement of the glass scale **(g)** showing strong angular dependence of the spectral reflectance. See also Figure 6.4.2 to illustrate the different 1D periodic structures.

(iii) To develop a semi-analytical (or even closed-form if possible) model that is predictive enough to avoid the need of a full EM simulation every time an optical parameter is changed (e.g., thicknesses, absorption, ...)

(iv) To infer an inverse model from the direct model obtained in (iii). This step is most of the time non-trivial because there are no analytical solutions. However, we could benefit from Deep Learning approaches to guide us towards an initial solution. Since Deep Learning approaches require a lot of data to train the network, in our case, this translates to generate a BSDF for each set of parameters; the benefit of a predictive BSDF model developed in (iii) is thus paramount.

In the long term, I wish to develop models of ageing and degradation of materials (cf. [DRS08]) in order to predict the appearance of materials over time and also to virtually restore the **appearance of ancient objects**. For this purpose, the combination of experimental archaeology, optical measurements and modelling is promising. Experimental archaeology allows the reconstruction of the object as it was manufactured at the time and the comparison of the appearance measurements of the original object with the restored object. These measurements can then be used to develop appearance models for ageing or restoration. Furthermore, the same models can then be applied to objects of the same period in order to restore them virtually, thus avoiding long measurement campaigns. Our recent work [31] is a first step in that direction.

In an **industrial context**, with the help of collaborations with material science experts, comparing measurements of artificially accelerated aged and new materials will allow the development of more predictive models than those usually developed in computer graphics (see [MG08] for an overview).

List of external references

[ACGO18] Rachel A. Albert, Dorian Yao Chan, Dan B. Goldman, and James F. O'Brien.
Approximate svBRDF Estimation From Mobile Phone Video.
In Wenzel Jakob and Toshiya Hachisuka, editors, *Eurographics Symposium on Rendering - Experimental Ideas & Implementations*. The Eurographics Association, 2018.

[APS00] Michael Ashikhmin, Simon Premoze, and Peter Shirley.
A microfacet-based BRDF generator.
In *Proc. ACM SIGGRAPH '00*, pages 65–74, 2000.

[AWL15] Miika Aittala, Tim Weyrich, and Jaakko Lehtinen.
Two-shot SVBRDF capture for stationary materials.
ACM Trans. Graph. (Proc. SIGGRAPH), 34(4):110:1–110:13, July 2015.

[BAEDR08] Aner Ben-Artzi, Kevin Egan, Fredo Durand, and Ravi Ramamoorthi.
A precomputed polynomial representation for interactive BRDF editing with global illumination.
ACM Trans. Graph., 27 (2):13 (1–13), May 2008.

[Bat21] Mégane Bati.
Design inverse de l'apparence de matériaux multicouches.
PhD thesis, 2021.
2021BORD0396.

[Bay21] Mélanie Marthe Bay.
The interplay between order and disorder in cholesteric hydroxypropyl cellulose films.
PhD thesis, University of Cambridge. Department of Chemistry, 2021.

[BJ19] Benedikt Bitterli and Wojciech Jarosz.
Selectively metropolised monte carlo light transport simulation.
ACM Trans. Graph., 38(6), nov 2019.

[Bli77] James F. Blinn.
Models of light reflection for computer synthesized pictures.

In *ACM SIGGRAPH '77*, pages 192–198, 1977.

[BLPW14] Adam Brady, Jason Lawrence, Pieter Peers, and Westley Weimer.
Genbrdf: Discovering new analytic brdfs with genetic programming.
ACM Trans. Graph., 33(4), jul 2014.

[BNM15] Samuel D. Butler, Stephen E. Nauyoks, and Michael A. Marciniak.
Experimental analysis of bidirectional reflectance distribution function cross section conversion term in direction cosine space.
Opt. Lett., 40(11):2445–2448, Jun 2015.

[BOK11] Ahmet Bilgili, Aydn Öztürk, and Murat Kurt.
A General BRDF Representation Based on Tensor Decomposition.
Computer Graphics Forum, 2011.

[BP20] James C. Bieron and Pieter Peers.
An adaptive brdf fitting metric.
Computer Graphics Forum, 39(4), July 2020.

[BS87] P. Beckmann and A. Spizzichino.
The scattering of electromagnetic waves from rough surfaces.
Artech House, 2 edition, 1987.

[BSH12] Mahdi M. Bagher, Cyril Soler, and Nicolas Holzschuch.
Accurate fitting of measured reflectances using a Shifted Gamma micro-facet distribution.
Computer Graphics Forum, 31(4):1509–1518, June 2012.

[BSN16] Mahdi M. Bagher, John Snyder, and Derek Nowrouzezahrai.
A non-parametric factor microfacet model for isotropic brdfs.
ACM Trans. Graph., 35(5), July 2016.

[Bur12] Brent Burley.
Physically-based shading at Disney.
In Stephen Hill and Stephen McAuley, editors, *Siggraph course: Practical Physically Based Shading in Film and Game Production*. ACM, 2012.

[BWP+20] Benedikt Bitterli, Chris Wyman, Matt Pharr, Peter Shirley, Aaron Lefohn, and Wojciech Jarosz.
Spatiotemporal reservoir resampling for real-time ray tracing with dynamic direct lighting.
ACM Trans. Graph., 39(4), aug 2020.

[CCFM22] O. Clausen, Y. Chen, A. Fuhrmann, and R. Marroquim.
Investigation and simulation of diffraction on rough surfaces.
Computer Graphics Forum, n/a(n/a), 2022.

[CDPS09] Massimiliano Corsini, Matteo Dellepiane, Federico Ponchio, and Roberto Scopigno.
Image-to-geometry registration: A Mutual Information method exploiting illumination-related geometric properties.
Computer Graphics Forum, 28(7):1755–1764, oct 2009.

[Cha60] S. Chandrasekhar.
Radiative Transfer.
Dover Books on Intermediate and Advanced Mathematics. Dover Publications, 1960.

[CJZ77] E. L. Church, H. A. Jenkinson, and J. M. Zavada.
Measurement of the Finish of Diamond-Turned Metal Surfaces By Differential Light Scattering.
Optical Engineering, 16(4):164360, 1977.

[CJZ79] E. L. Church, H. A. Jenkinson, and J. M. Zavada.
Relationship between Surface Scattering and Microtopographic Features.
Optical Engineering, 18(2):182125, 1979.

[CKS+17] Chakravarty R. Alla Chaitanya, Anton S. Kaplanyan, Christoph Schied, Marco Salvi, Aaron Lefohn, Derek Nowrouzezahrai, and Timo Aila.
Interactive reconstruction of monte carlo image sequences using a recurrent denoising autoencoder.
ACM Trans. Graph., 36(4), July 2017.

[CSCC14] David Coleman, Ioan Sucan, Sachin Chitta, and Nikolaus Correll.
Reducing the barrier to entry of complex robotic software: a moveit! case study, 2014.

[CT82] R. L. Cook and Kenneth E. Torrance.
A reflectance model for computer graphics.
ACM Trans. Graph., 1(1):7–24, 1982.

[CT86] E L. Church and P. Z. Takacs.
Statistical and signal processing concepts in surface metrology.
In *Proc. SPIE 0645, Optical Manufacturing, Testing and Aspheric Optics*, volume 0645, pages 107–115, 1986.

[dGM07] Luca di Gaspero and Eric Moyer.
Quadprog++.
https://sourceforge.net/projects/quadprog/, 2007.

[dGM16] Luca di Gaspero and Eric Moyer.
Quadprog++.
https://github.com/liuq/QuadProgpp, 2016.

[DGND08] Joannopoulos John D., Johnson Steven G., Winn Joshua N., and Meade Robert D.
Photonic Crystals, Molding the Flow of Light.
Princeton University Press, 2008.

[DRS08] Julie Dorsey, Holly Rushmeier, and Franois Sillion.
Digital Modeling of Material Appearance.
Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2008.

[DSHL10] Holger Dammertz, Daniel Sewtz, Johannes Hanika, and Hendrik P. A. Lensch.
Edge-avoiding À-trous wavelet transform for fast global illumination filtering.
In *Proceedings of the Conference on High Performance Graphics*, HPG '10, page 67–75, Goslar, DEU, 2010. Eurographics Association.

[DWT+10] Yue Dong, Jiaping Wang, Xin Tong, John Snyder, Yanxiang Lan, Moshe Ben-Ezra, and Baining Guo.
Manifold bootstrapping for SVBRDF capture.
ACM Trans. Graph., 29:98 (1–10), July 2010.

[ECT+21] Amos Egel, Krzysztof M. Czajkowski, Dominik Theobald, Konstantin Ladutenko, Alexey S. Kuznetsov, and Lorenzo Pattelli.
Smuthi: A python package for the simulation of light scattering by multiple particles near or between planar interfaces.
Journal of Quantitative Spectroscopy and Radiative Transfer, 273:107846, 2021.

[EG04] Tanos Mikhael Elfouhaily and Charles-Antoine Guérin.
A critical survey of approximate scattering wave theories from random rough surfaces.
Waves in Random Media, 14(4):R1–R40, 2004.

[FHF+17] Luca Fascione, Johannes Hanika, Marcos Fajardo, Per Christensen, Brent Burley, and Brian Green.
Path tracing in production - part 1: Production renderers.
In *ACM SIGGRAPH 2017 Courses*, SIGGRAPH '17, pages 13:1–13:39, 2017.

[FJM+20] J. R. Frisvad, S. A. Jensen, J. S. Madsen, A. Correia, L. Yang, S. K. S. Gregersen, Y. Meuret, and P.-E. Hansen.
Survey of models for acquiring the optical properties of translucent materials.
Computer Graphics Forum, 39(2):729–755, 2020.

[FVH+13] J. Filip, R. Vávra, M. Haindl, P. id, M. Krupika, and V. Havran.
Brdf slices: Accurate adaptive anisotropic appearance acquisition.
In *2013 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1468–1473, June 2013.

[Gal19] Alain Galvan.
Machine learning denoising. a review of different techniques for denoising monte-carlo images with machine learning.
https://alain.xyz/blog/machine-learning-denoising, Nov. 2019.

[GGG+16] D. Guarnera, G. C. Guarnera, A. Ghosh, C. Denk, and M. Glencross.
BRDF representation and acquisition.
Computer Graphics Forum, 35(2):625–650, may 2016.

[GGSK19] P. Grittmann, I. Georgiev, P. Slusallek, and J. Křivánek.
Variance-aware multiple importance sampling.
ACM Trans. Graph. (SIGGRAPH Asia 2019), 38(6), 2019.

[GJZ21] Yu Guo, Adrian Jarabo, and Shuang Zhao.
Beyond mie theory: Systematic computation of bulk scattering parameters based on microphysical wave optics.
ACM Trans. Graph., 40(6), dec 2021.

[GKPB04] Pascal Gautron, Jaroslav Křivánek, Sumanta N. Pattanaik, and Kadi Bouatouch.
A novel hemispherical basis for accurate and efficient rendering.
In Henrik Wann Jensen and Alexander Keller, editors, *Proc. Eurographics Symposium on Rendering 2004*, pages 321–330. Eurographics Association, 2004.

[GKPS12] I. Georgiev, J. Křivánek, S. Popov, and P. Slusallek.
Importance caching for complex illumination.
In *Computer Graphics Forum*, volume 31, pages 701–710. Wiley Online Library, 2012.

[Gre09] Eugène Grenon.
Observer l'extrêmement petit: au microscope, c'est possible, 2009.

[Gue20] G. Guennebaud.
Eigen version of the Quadprog++ library.
https://github.com/ggael/QuadProg, 2020.

[Hal07] Johan Hallenberg.
Robot Tool Center Point Calibration using Computer Vision Master ' s Thesis in Computer Vision by.

Master's thesis, Linköping University, 2007.

[Har75] James E. Harvey.
Light-Scattering Characteristics of Optical Surface.
PhD thesis, University of Arizona, 1975.
Adviser: R. V. Shack.

[Hei14] Eric Heitz.
Understanding the Masking-Shadowing Function in Microfacet-Based BRDFs.
Journal of Computer Graphics Techniques, 3(2):32–91, June 2014.

[Hes95] Tim Hesterberg.
Weighted Average Importance Sampling and Defensive Mixture Distributions.
Technometrics, 37(2):185–194, 1995.

[HEV+16a] S. Herholz, O. Elek, J. Vorba, H. Lensch, and J. Křivánek.
Product importance sampling for light transport path guiding.
In *Computer Graphics Forum*, volume 35, pages 67–77. Wiley Online Library, 2016.

[HEV+16b] Sebastian Herholz, Oskar Elek, Jiří Vorba, Hendrik Lensch, and Jaroslav Křivánek.
Product importance sampling for light transport path guiding.
Computer Graphics Forum, 35(4):67–77, 2016.

[HHN+17] Vlastimil Havran, Jan Hošek, Šárka Němcová, Jiří Čáp, and Jiří Bittner.
Lightdrum—Portable light stage for accurate BTF measurement on site.
Sensors (Switzerland), 17(3), 2017.

[HJBS79] B. J. Hoenders, E. Jakeman, H. P. Baltes, and B . Steinle.
K correlations and facet models in diffuse scattering.
Optica Acta: International Journal of Optics, 26(10):1307–1319, 1979.

[HLZ10] Michael Holroyd, Jason Lawrence, and Todd Zickler.
A coaxial optical scanner for synchronous acquisition of 3D geometry and surface reflectance.
ACM SIGGRAPH 2010 Papers, SIGGRAPH 2010, 29(4):1, jul 2010.

[HP03] Jefferson Y. Han and Ken Perlin.
Measuring bidirectional texture reflectance with a kaleidoscope.
ACM Trans. Graph., 22(3):741–748, July 2003.

[HS13] Jakub Haberko and Frank Scheffold.
Fabrication of mesoscale polymeric templates for three-dimensional disordered photonic materials.
Opt. Express, 21(1):1057–1065, Jan 2013.

[HSCD12] James E. Harvey, Sven Schröder, Narak Choi, and Angela Duparré.
Total integrated scatter from surfaces with arbitrary roughness, correlation widths, and incident angles.
Optical Engineering, 51(1):013402–1–013402–11, 2012.

[HST10] Kaiming He, Jian Sun, and Xiaoou Tang.
Guided image filtering.
In Kostas Daniilidis, Petros Maragos, and Nikos Paragios, editors, *Computer Vision – ECCV 2010*, pages 1–14, Berlin, Heidelberg, 2010. Springer Berlin Heidelberg.

[HTSG91] Xiao D. He, Kenneth E. Torrance, François X. Sillion, and Donald P. Greenberg.
A comprehensive physical model for light reflection.
In *Proceedings of the 18th Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH '91, page 175–186, New York, NY, USA, 1991. Association for Computing Machinery.

[Int19] Intel open image denoise. high performance denoising library for ray-tracing.
https://openimagedenoise.github.io/, 2019.

[JEK+19] Wojciech Jarosz, Afnan Enayet, Andrew Kensler, Charlie Kilpatrick, and Per Christensen.
Orthogonal array sampling for Monte Carlo rendering.
Computer Graphics Forum (Proceedings of EGSR), 38(4):135–147, July 2019.

[Jen01] Henrik Wann Jensen.
Realistic Image Synthesis Using Photon Mapping, page 208.
A.K. Series / CRC Press, 2001.

[JMLH01] H. W. Jensen, S. R. Marschner, M. Levoy, and P. Hanrahan.
A practical model for subsurface light transport.
In *Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH '01, page 511–518. Association for Computing Machinery, 2001.

[JS06] Nocedal J. and Wright S.
Numerical Optimization.
Springer-Verlag, New York Inc., July 2006.

[Kaj86] James T. Kajiya.
The rendering equation.
SIGGRAPH Comput. Graph., 20(4):143–150, August 1986.

[KBS15] Nima Khademi Kalantari, Steve Bako, and Pradeep Sen.
A machine learning approach for filtering monte carlo noise.
ACM Trans. Graph., 34(4), July 2015.

[KDS96] Jan Koenderink, Andrea Doorn, and Marigo Stavridi.
BRDF expressed in terms of surface scattering modes.
In *Proc. Eur. Conf. on Comp. Vision '96*, pages 28–39, 1996.

[KNRS13] Johannes Kohler, Tobias Noll, Gerd Reis, and Didier Stricker.
A full-spherical device for simultaneous geometry and reflectance acquisition.
In *Proceedings of IEEE Workshop on Applications of Computer Vision*, pages 355–362, 2013.

[KoCF06] A. Krywonos and University of Central Florida.
Predicting surface scatter using a linear systems formulation of non-paraxial scalar diffraction, 2006.

[Kry06] Andrey Krywonos.
Predicting Surface Scatter using a Linear Systems Formulation of Non-Paraxial Scalar Diffraction.
PhD thesis, University of Central Florida, 2006.
Adviser: J. E. Harvey.

[KVG+19] I. Kondapaneni, P. Vévoda, P. Grittmann, T. Skřivan, P. Slusallek, and J. Křivánek.
Optimal multiple importance sampling.
ACM Trans. Graph., 38(4), 2019.

[KW09] Malvin H. Kalos and Paula A. Whitlock.
Monte Carlo Methods.
John Wiley & Sons, Ltd, 2009.

[LBGO15] Ronan Le Breton, Guillaume Ged, and Gaël Obein.
Out of plane brdf measurement at lne-cnam using "condor" our primary goniospectrophotometer.
In *28th CIE SESSION, Manchester, United Kingdom*, volume 1, pages 1401–1407. CIE 216:2015, 2015.

[Len89] Reimar K. Lenz.
A New Technique for Fully Autonomous and Efficient 3D Robotics Hand/Eye Calibration.
IEEE Transactions on Robotics and Automation, 5(3):345–358, 1989.

[LF12] Victor Liu and Shanhui Fan.
S^4 : A free electromagnetic solver for layered periodic structures.
Computer Physics Communications, 183(10):2233 – 2244, 2012.

[LFTG97] Eric P. F. Lafortune, Sing-Choong Foo, Kenneth E. Torrance, and Donald P. Greenberg.
Non-linear approximation of reflectance functions.
In *ACM SIGGRAPH '97*, pages 117–126, 1997.

[LKB+22] Daqi Lin, Markus Kettunen, Benedikt Bitterli, Jacopo Pantaleoni, Cem Yuksel, and Chris Wyman.
Generalized resampled importance sampling: Foundations of restir.
ACM Trans. Graph., 41(4), jul 2022.

[LKYU12] Joakim Löw, Joel Kronander, Anders Ynnerman, and Jonas Unger.
Brdf models for accurate and efficient rendering of glossy surfaces.
ACM Trans. Graph., 31(1), February 2012.

[LLCK19] Edward Liu, Ignacio Llamas, Juan Cañada, and Patrick Kelly.
Cinematic rendering in ue4 with real-time ray tracing and denoising.
In Eric Haines and Tomas Akenine-Möller, editor, *Ray Tracing Gems*, chapter 19, pages 289–319. Apress, 2019.

[Lou04] M. I. A. Lourakis.
levmar: Levenberg-marquardt nonlinear least squares algorithms in c/c++, 2004.

[LRR04] Jason Lawrence, Szymon Rusinkiewicz, and Ravi Ramamoorthi.
Efficient brdf importance sampling using a factored representation.
In *ACM SIGGRAPH 2004 Papers*, SIGGRAPH '04, page 496–505, New York, NY, USA, 2004. Association for Computing Machinery.

[Lu14] Heqi Lu.
Importance Sampling of Realistic Light Sources.
Theses, Université de Bordeaux, February 2014.

[Luc20] Antoine Lucat.
Acquisition opto-numérique de vêtements asiatiques anciens.
PhD thesis, Université de Bordeaux, 2020.
2020BORD0012.

[MAA01] Michael D. McCool, Jason Ang, and Anis Ahmad.
Homomorphic factorization of brdfs for high-performance rendering.

In *Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH '01, page 171–178, New York, NY, USA, 2001. Association for Computing Machinery.

[Mar98] Stephen R Marschner.
Inverse rendering for computer graphics.
PhD thesis, Cornell University, 1998.

[Mat22] Mathworks.
Matlab implementation of constrained nonlinear optimization algorithms, 2022.

[McA02] David K. McAllister.
A generalized surface appearance representation for computer graphics.
thesis, Chapel Hill, page 129, 2002.

[Mes88] J.B. Messerschmitt.
Ueber diffuse reflexion.
Annalen der Physik, page 867, 1888.

[MG81] M. G. Moharam and T. K. Gaylord.
Rigorous coupled-wave analysis of planar-grating diffraction.
J. Opt. Soc. Am., 71(7):811–818, Jul 1981.

[MG08] S. Mérillou and D. Ghazanfarpour.
Technical section: A survey of aging and weathering phenomena in computer graphics.
Comput. Graph., 32(2):159–174, April 2008.

[MMBJ17] Michael Mara, Morgan McGuire, Benedikt Bitterli, and Wojciech Jarosz.
An efficient denoising algorithm for global illumination.
In *Proceedings of High Performance Graphics*, New York, NY, USA, July 2017. ACM.

[MMS+04] Gero Müller, Jan Meseth, Mirko Sattler, Ralf Sarlette, and Reinhard Klein.
Acquisition, synthesis and rendering of bidirectional texture functions.
In Christophe Schlick and Werner Purgathofer, editors, *Eurographics 2004, State of the Art Reports*, pages 69–94. INRIA and Eurographics Association, September 2004.

[MPBM03] Wojciech Matusik, Hanspeter Pfister, Matt Brand, and Leonard McMillan.
A data-driven reflectance model.
ACM Trans. Graph., 22(3):759–769, July 2003.

[MTM96] Michael I. Mishchenko, Larry D. Travis, and Daniel W. Mackowski.

T-matrix computations of light scattering by nonspherical particles: A review.
J. Quant. Spectrosc. Radiat. Transf., 55(5):535–575, 1996.

[NB11] Colette Naud and André Bergeron.
La lumière et l'éclairage, 2011.

[NDM05] A. Ngan, F. Durand, and W. Matusik.
Experimental Analysis of BRDF Models.
In *Proc. Eurographics Symposium on Rendering*, pages 117–226, 2005.

[NKRS15] Tobias Nöll, Johannes Köhler, Gerd Reis, and Didier Stricker.
Fully automatic, omnidirectional acquisition of geometry and appearance in the context of cultural heritage preservation.
Journal on Computing and Cultural Heritage, 8:1–28, 02 2015.

[NRH+77] F. E. Nicodemus, J. C. Richmond, J. J. Hsia, I. W. Ginsberg, and T. Limperis.
Geometrical Considerations and Nomenclature for Reflectance.
National Bureau of Standards, 1977.

[NVi18] NVidia.
Nvidia rtx platform.
https://developer.nvidia.com/rtx, 2018.

[OLK+21] Y. Ouyang, S. Liu, M. Kettunen, M. Pharr, and J. Pantaleoni.
Restir gi: Path resampling for real-time path tracing.
Computer Graphics Forum, 40(8):17–29, 2021.

[Opt19] Nvidia optix ai-accelerated denoiser.
https://developer.nvidia.com/optix-denoiser, 2019.

[PBD+10] Steven G. Parker, James Bigler, Andreas Dietrich, Heiko Friedrich, Jared Hoberock, David Luebke, David McAllister, Morgan McGuire, Keith Morley, Austin Robison, and Martin Stich.
Optix: A general purpose ray tracing engine.
ACM Trans. Graph., 29(4), jul 2010.

[PBPP11] Anthony Pajot, Loïc Barthe, Mathias Paulin, and Pierre Poulin.
Representativity for robust and adaptive multiple importance sampling.
IEEE Trans. Visualization and Computer Graphics, 17(8):1108–1121, 2011.

[PG10] J.M. Palmer and B.G. Grant.
The Art of Radiometry.
Press Monographs. SPIE Press, 2010.

[PH10] Matt Pharr and Greg Humphreys.
Physically Based Rendering: From Theory to Implementation.
Morgan Kaufmann, 2 edition, 2010.

[Pho75] Bui Tuong Phong.
Illumination for computer generated pictures.
Commun. ACM, 18(6):311–317, June 1975.

[PJH16] Matt Pharr, Wenzel Jakob, and Greg Humphreys.
Physically Based Rendering.
Morgan Kaufmann, November 2016.

[Ray07] Lord Rayleigh.
On the dynamical theory of gratings.
Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character, 79(532):399–416, 1907.

[Rei08] Vincent Reillon.
Characterization and modelling of the optical and colorimetric properties of lustred ceramics.
Theses, Université Pierre et Marie Curie - Paris VI, October 2008.

[Ric51a] Stephen O. Rice.
Reflection of electromagnetic waves from slightly rough surfaces.
Communications on Pure and Applied Mathematics, 4(2-3):351–378, 1951.

[Ric51b] Stephen O. Rice.
Reflection of electromagnetic waves from slightly rough surfaces.
Communications on Pure and Applied Mathematics, 4(2-3):351–378, 1951.

[Ric72] William Hadley Richardson.
Bayesian-based iterative method of image restoration∗.
J. Opt. Soc. Am., 62(1):55–59, Jan 1972.

[Rus98] Szymon Rusinkiewicz.
A new change of variables for efficient BRDF representation.
In *Proc. EGWR '98*, pages 11–22, 1998.

[RVZ08] Fabiano Romeiro, Yuriy Vasilyev, and Todd Zickler.
Passive reflectometry.
In *Proc. Eur. Conf. on Comp. Vision*, pages 859–872, 2008.

[RWS+11] Peiran Ren, Jiaping Wang, John Snyder, Xin Tong, and Baining Guo.

Pocket reflectometry.
In *ACM SIGGRAPH 2011 Papers*, SIGGRAPH '11, pages 45:1–45:10, New York, NY, USA, 2011. ACM.

[SAWG91] Françis X. Sillion, James R. Arvo, Stephen H. Westin, and Donald P. Greenberg.
A global illumination solution for general reflectance distributions.
In *Proc. ACM SIGGRAPH '91*, pages 187–196, 1991.

[SBN15] Cyril Soler, Mahdi M. Bagher, and Derek Nowrouzezahrai.
Efficient and accurate spherical kernel integrals using isotropic decomposition.
ACM Trans. Graph., 34(5), November 2015.

[SCV07] O. Salazar Celis, A. Cuyt, and B. Verdonk.
Rational approximation of vertical segments.
Numerical Algorithms, 45:375–388, 2007.

[SDC+11] Sven Schröder, Angela Duparré, Luisa Coriand, Andreas Tünnermann, Dayana H. Penalver, and James E. Harvey.
Modeling of light scattering in different regimes of surface roughness.
Opt. Express, 19(10):9820–9835, May 2011.

[SHSK19] M. Sbert, V. Havran, and L. Szirmay-Kalos.
Optimal deterministic mixture sampling.
In *Eurographics (Short Papers)*, pages 73–76, 2019.

[Smi67] B. Smith.
Geometrical shadowing of a random rough surface.
IEEE Transactions on Antennas and Propagation, 15(5), September 1967.

[Sol21] Cyril Soler.
private communication, February 2021.

[SP94] François X. Sillion and Claude Puech.
Radiosity and Global Illumination, volume 11.
Morgan Kaufmann Publishers, 1994.

[SR09] Olga Sorkine and Michael Rabinovich.
Least-squares rigid motion using svd.
Technical Report February, 2009.

[SSK03] Mirko Sattler, Ralf Sarlette, and Reinhard Klein.
Efficient and realistic visualization of cloth.

In *Proceedings of the 14th Eurographics Workshop on Rendering*, EGRW '03, pages 167–177, Aire-la-Ville, Switzerland, Switzerland, 2003. Eurographics Association.

[SSW+14] Christopher Schwartz, Ralf Sarlette, Michael Weinmann, Martin Rump, and Reinhard Klein.
Design and implementation of practical bidirectional texture function measurement devices focusing on the developments at the university of bonn.
Sensors, 14(5), April 2014.

[SSWK13a] Christopher Schwartz, Ralf Sarlette, Michael Weinmann, and Reinhard Klein.
Dome ii: A parallelized btf acquisition system.
In Holly Rushmeier and Reinhard Klein, editors, *Eurographics Workshop on Material Appearance Modeling: Issues and Acquisition*, pages 25–31. Eurographics Association, June 2013.

[SSWK13b] Christopher Schwartz, Ralf Sarlette, Michael Weinmann, and Reinhard Klein.
DOME II: a parallelized BTF acquisition system.
Workshop on Material Appearance Modeling (2013), pages 25–31, 2013.

[Sta99] Jos Stam.
Diffraction shaders.
In *SIGGRAPH '99*. ACM, 1999.

[Sta14] Doekele Stavenga.
Thin film and multilayer optics cause structural colors of many insects and birds.
Materials Today: Proceedings, 1, 12 2014.

[Sto95] John C. Stover.
Optical Scattering: Measurement and Analysis.
SPIE Optical Engineering Press, 1995.

[SWR+11] Christopher Schwartz, Michael Weinmann, Roland Ruiters, Arno Zinke, Ralf Sarlette, and Reinhard Klein.
Capturing shape and reflectance of food.
In *SIGGRAPH Asia 2011 Sketches*, SA '11, pages 28:1–28:2, New York, NY, USA, December 2011. ACM.

[SWRK11] Christopher Schwartz, Michael Weinmann, Roland Ruiters, and Reinhard Klein.
Integrated High-Quality Acquisition of Geometry and Appearance for Cultural Heritage.
In *The 12th International Symposium on Virtual Reality, Archeology and Cultural Heritage VAST 2011*, pages 25–32, 2011.

[TCE05] Justin F. Talbot, David Cline, and Parris Egbert.
Importance resampling for global illumination.
In *Proceedings of the Sixteenth Eurographics Conference on Rendering Techniques*, EGSR '05, page 139–146, Goslar, DEU, 2005. Eurographics Association.

[Tha03] Fritz Thaler.
Die diffuse reflexion des lichtes an matten oberflächen.
Annalen der Physik, 316(8):996–1019, 1903.

[TS66] K. E. Torrance and E. M. Sparrow.
Off-Specular Peaks in the Directional Distribution of Reflected Thermal Radiation.
Journal of Heat Transfer, 88(2):223–230, 05 1966.

[TU+20] Tanaboon Tongbuasirilai, Jonas Unger, , Joel Kronander, and Murat Kurt.
Compact and intuitive data-driven brdf models.
The Visual Computer, 36(4):855–872, Apr 2020.

[Vea97] Eric Veach.
Robust Monte Carlo methods for light transport simulation, volume 1610.
Stanford University PhD thesis, 1997.

[Vea98a] Eric Veach.
Robust Monte Carlo Methods for Light Transport Simulation.
PhD thesis, Stanford University, Stanford, CA, USA, 1998.
AAI9837162.

[Vea98b] Eric Veach.
Robust Monte Carlo Methods for Light Transport Simulation.
PhD thesis, Stanford University, Stanford, CA, USA, 1998.
AAI9837162.

[VH98] C. L. Vernold and J. E. Harvey.
A modified beckmann-kirchhoff scattering theory.
Proc. SPIE, 3426, Scattering and Surface Roughness II:51–56, 1998.

[WAT92] Stephen Westin, James Arvo, and Kenneth Torrance.
Predicting reflectance functions from complex surfaces.
In *Proc. ACM SIGGRAPH '92*, pages 255–264, 1992.

[WDR11] Hongzhi Wu, Julie Dorsey, and Holly Rushmeier.
Physically-based interactive bi-scale material design.
In *Proceedings of the 2011 SIGGRAPH Asia Conference*, SA '11, New York, NY, USA, 2011. Association for Computing Machinery.

[WK15] Michael Weinmann and Reinhard Kleiny.
Advances in geometry and reflectance acquisition (Course Notes).
In *SIGGRAPH Asia 2015 Courses, SA 2015*. Association for Computing Machinery, Inc, nov 2015.

[WMLT07] B. Walter, S. Marschner, H. Li, and K. E. Torrance.
Microfacet models for refraction through rough surfaces.
In *Eurographics Symposium on Rendering*, pages 195–206, 2007.

[WND+14] A. Wilkie, S. Nawaz, M. Droske, A. Weidlich, and J. Hanika.
Hero wavelength spectral sampling.
In *Proceedings of the 25th Eurographics Symposium on Rendering*, EGSR '14, pages 123–131, Aire-la-Ville, Switzerland, Switzerland, 2014. Eurographics Association.

[WW07] A. Weidlich and A. Wilkie.
Arbitrarily layered micro-facet surfaces.
In *Proc. of the 5th International Conference on Computer Graphics and Interactive Techniques in Australia and Southeast Asia*, GRAPHITE '07, page 171–178. ACM, 2007.

[WW09] A. Weidlich and A. Wilkie.
Exploring the potential of layered brdf models.
In *ACM SIGGRAPH ASIA 2009 Courses*, SIGGRAPH ASIA '09, pages 7:1–7:58. ACM, 2009.

[Yee66] Kane S. Yee.
Numerical solution of initial boundary value problems involving maxwell's equations in isotropic media.
IEEE Trans. Antennas and Propagation, pages 302–307, 1966.

[YHW+18] Ling-Qi Yan, Miloš Hašan, Bruce Walter, Steve Marschner, and Ravi Ramamoorthi.
Rendering specular microgeometry with wave optics.
ACM Trans. Graph., 37(4), jul 2018.

[YLL+18] Bo Yang, Wenwei Liu, Zhancheng Li, Hua Cheng, Shuqi Chen, and Jianguo Tian.
Polarization-sensitive structural colors with hue-and-saturation tuning based on all-dielectric nanopixels.
Advanced Optical Materials, 6(4):1701009, 2018.

[ZCD+16] Zhiming Zhou, Guojun Chen, Yue Dong, David Wipf, Yong Yu, John Snyder, and Xin Tong.

Sparse-as-possible svbrdf acquisition.
ACM Trans. Graph., 35(6):189:1–189:12, November 2016.

[ZERB05] Todd Zickler, Sebastian Enrique, Ravi Ramamoorthi, and Peter Belhumeur.
Reflectance sharing: Image-based rendering from a sparse set of images.
In *Proc. Eurographics Symposium on Rendering*, pages 253–265, 2005.

[Zha00] Zhengyou Zhang.
A flexible new technique for camera calibration.
Technical Report 11, Microsoft Research, 2000.

[ZJL+15] M. Zwicker, W. Jarosz, J. Lehtinen, B. Moon, R. Ramamoorthi, F. Rousselle, P. Sen, C. Soler, and S.-E. Yoon.
Recent advances in adaptive sampling and reconstruction for monte carlo rendering.
Comput. Graph. Forum, 34(2):667–681, May 2015.

[ZTZ05] O. C. Zienkiewicz, R. L. Taylor, and J. Z. Zhu.
The Finite Element Method: Its Basis and Fundamentals, Sixth Edition.
Butterworth-Heinemann, 6 edition, May 2005.

BIBLIOGRAPHY

[1] Ramon Hegedus, Antoine Lucat, Justine Redon, and Romain Pacanowski.
Isotropic BRDF Measurements with Quantified Uncertainties.
In *EUROGRAPHICS WORKSHOP ON MATERIAL APPEARANCE MODELING*, Dublin, Ireland, June 2016.

[2] Antoine Lucat, Ramon Hegedus, and Romain Pacanowski.
Diffraction Removal in an Image-based BRDF Measurement Setup.
In *EI 2018 - Electronic Imaging Material Appearance 2018* , pages 1–6, Burlingame, California, United States, January 2018.

[3] Antoine Lucat, Ramon Hegedus, and Romain Pacanowski.
Diffraction effects detection for HDR image-based measurements.
Optics Express, 25(22):2921 – 2929, October 2017.

[4] Antoine Lucat, Ramon Hegedus, and Romain Pacanowski.
Diffraction Prediction in HDR measurements.
In *EUROGRAPHICS WORKSHOP ON MATERIAL APPEARANCE MODELING*, Helsinki, Finland, 2017.

[5] Romain Pacanowski, Xavier Granier, Christophe Schlick, and Poulin Pierre.
Sketch and Paint-based Interfacefor Highlight Modeling.
In *Fifth Eurographics conference on Sketch-Based Interfaces and Modeling*, Annecy, France, June 2008.

[6] Romain Vergne, Romain Pacanowski, Pascal Barla, Xavier Granier, and Christophe Schlick.
Light Warping for Enhanced Surface Depiction.
ACM Transactions on Graphics, 28(3):25:1–25:8, July 2009.

[7] Romain Vergne, Romain Pacanowski, Pascal Barla, Xavier Granier, and Christophe Schlick.
Improving Shape Depiction under Arbitrary Rendering.
IEEE Transactions on Visualization and Computer Graphics, 17(8):1071 – 1081, June 2011.

[8] Boris Raymond, Pascal Barla, Romain Pacanowski, and Xavier Granier.
Understanding Anisotropic Highlights for Interactive Appearance Editing.

PRISM2: The science of light & shade, October 2013.
Poster.

[9] Romain Pacanowski, Oliver Salazar-Celis, Christophe Schlick, Xavier Granier, Poulin Pierre, and Cuyt Annie.
Rational BRDF.
IEEE Transactions on Visualization and Computer Graphics, 18(11):1824–1835, March 2012.

[10] Laurent Belcour, Romain Pacanowski, Marion Delahaie, Aude Laville-Geay, and Laure Eupherte.
BRDF Measurements and Analysis of Retroreflective Materials.
Journal of the Optical Society of America, 31(12):2561–2572, December 2014.

[11] Kevin Vynck, Romain Pacanowski, Pascal Barla, Xavier Granier, and Philippe Lalanne.
Engineering nanostructured surfaces with correlated disorder for visual appearance design.
In *Workshop on "Correlated Disorder and hyperuniformity in soft matter and photonics"*, Paris, France, October 2018.

[12] Heqi Lu, Romain Pacanowski, and Xavier Granier.
Position-Dependent Importance Sampling of Light Field Luminaires.
IEEE Transactions on Visualization and Computer Graphics, 21(2):241 – 251, 2015.

[13] Romain Pacanowski, Xavier Granier, Christophe Schlick, and Poulin Pierre.
Volumetric Vector-based Representation for Indirect Illumination Caching.
Journal of Computer Science and Technology (JCST), 25(5):925–932, September 2010.

[14] Heqi Lu, Romain Pacanowski, and Xavier Granier.
Real-Time Importance Sampling of Dynamic Environment Maps.
In *Eurographics 2013 - Short Papers*, pages 65–68, Girona, Spain, May 2013. Eurographics.

[15] Heqi Lu, Romain Pacanowski, and Xavier Granier.
Second-Order Approximation for Variance Reduction in Multiple Importance Sampling.
Computer Graphics Forum, 32(7):131–136, October 2013.

[16] David Murray, Sofiane Benzait, Romain Pacanowski, and Xavier Granier.
On Learning the Best Balancing Strategy.
In *Eurographics 2020 - Annual Conference of the European Association for Computer Graphics*, volume xx, pages 1 – 4, Norrköping / Virtual, Sweden, May 2020.

[17] Arthur Dufay, Pascal Lecocq, Romain Pacanowski, Jean-Eudes Marvie, and Xavier Granier.
Cache-friendly micro-jittered sampling.
In *SIGGRAPH 2016*, Anaheim, United States, July 2016.

[18] Mégane Bati, Romain Pacanowski, and Pascal Barla.
Comparative study of layered material models.
In H. Rushmeier and R. Klein, editors, *Workshop on Material Appearance Modeling*, Strasbourg, France, July 2019.

[19] Mégane Bati, Pascal Barla, and Romain Pacanowski.
An Inverse Method for the Exploration of Layered Material Appearance.
ACM Transactions on Graphics, August 2021.

[20] Laurent Belcour, Oliver Salazar-Celis, and Romain Pacanowski.
An Addendum to Rational Function Fitting.
Technical report, p., December 2013.
This report is an addendum to the Rational BRDF publication available in hal-00678885.

[21] Laurent Belcour, Pascal Barla, and Romain Pacanowski.
ALTA: A BRDF analysis library.
In Holly Rushmeier and Reinhard Klein, editors, *EGSR Workshop on Material Appearance Modeling*, Workshop on Material Appearance Modeling 2014, Lyon, France, June 2014. Eurographics.

[22] Nicolas Holzschuch and Romain Pacanowski.
A Two-Scale Microfacet Reflectance Model Combining Reflection and Diffraction.
ACM Transactions on Graphics, 36(4):12, July 2017.
Article 66.

[23] Pascal Barla, Laurent Belcour, and Romain Pacanowski.
In Praise of an Alternative BRDF Parametrization.
In *Workshop on Material Appearance Modeling*, Proceedings of the Workshop on Material Appearance Modeling 2015, Darmstadt, Germany, June 2015.

[24] Romain Pacanowski.
I3D Presentation of TVCG Paper Rational BRDF.
In *ACM Siggraph Symposium on Interactive 3D Graphics and Games*, Orlando, United States, March 2013.

[25] Nicolas Holzschuch and Romain Pacanowski.
A physically accurate reflectance model combining reflection and diffraction.
Research Report RR-8807, INRIA, November 2015.

[26] Nicolas Holzschuch and Romain Pacanowski.
The Effects of Digital Cameras Optics and Electronics for Material Acquisition.
In H. Rushmeier and R. Klein, editors, *Workshop on Material Appearance Modeling*, pages 1–4, Helsinki, Finland, June 2017.

[27] Arthur Dufay, David Murray, Romain Pacanowski, et al.
The Malia Rendering Framework.
https://pacanows.gitlabpages.inria.fr/MRF/main.md.html, 2019.

[28] Mégane Bati, Romain Pacanowski, and Pascal Barla.
Numerical Analysis of Layered Materials Models.
Research report, Institut d'Optique Graduate School ; Université de Bordeaux ; CNRS ; INRIA, June 2019.

[29] David Murray, Alban Fichet, and Romain Pacanowski.
Efficient Spectral Rendering on the GPU for Predictive Rendering.
In *Ray Tracing Gems II*, pages 673 – 698. Springer, 2021.

[30] Kevin Vynck, Romain Pacanowski, Adrian Agreda, Arthur Dufay, Xavier Granier, and Philippe Lalanne.
The visual appearances of disordered optical metasurfaces.
Nature Materials, 21:1035–1041, May 2022.

[31] Aurélie Mounier, Markus Schlicht, Maud Mulliez, Romain Pacanowski, Antoine Lucat, and Pascal Mora.
In search of the lost polychromy of English medieval alabaster panels in the Southwest of France.
Color Research and Application, February 2020.