Loughborough University

LOUGHBOROUGH UNIVERSITY

The Department of Computer Science

Doctoral research in Computer Graphics

---

# IMAGE BASED SURFACE REFLECTANCE REMAPPING FOR CONSISTENT AND TOOL INDEPENDENT MATERIAL APPEARANCE

by DAR'YA GUARNERA

Doctoral Thesis submitted in partial fulfilment of the requirements
for the award of
Doctor of Philosophy of Loughborough University

September 2017

FIRST SUPERVISOR
Dr. Baihua Li

SECOND SUPERVISOR
Professor Eran Edirisinghe



---

# Contents

# List of Figures

# Acknowledgements

First and foremost I thank Dr. Giuseppe Claudio Guarnera with all my heart for all the support and guidance throughout my Ph.D. research: this work would not have happen without his support. I would also like to thank him for teaching me and providing me with the skills required for my project and all the help with experiments and publications. Claudios support was not only shown in professional activities but also in our personal relationships as my wonderful, loving, supportive, encouraging and patient husband, who has being with me through all my ups and downs.

I would like to sincerely thank to my former PhD supervisor Dr Mashhuda Glencross, for granting me the opportunity to engage in my Ph.D. research, for the support of my Ph.D. studies and related research and for her patience with all the proof reading. Her guidance helped me a lot with the research and even in writing this thesis.

I also would like to thank BMW representative for the opportunity to gain experience within the industry workflow. It was an interesting and useful experience to participate in collaborations between departments.

I thank my parents Nelly and Alexander Shaposhnikovi and my daughter Anastasiya for supporting me spiritually through my PhD. I would also like to thank all of my friends who supported and encouraged me to strive towards my goal.

Finally would like to thank my current supervisors Dr. Baihua Li and Prof. Eran Edirisinghe for their help, feedback and support.

# Abstract

Physically-based rendering in Computer Graphics requires the knowledge of material properties other than 3D shapes, textures and colors, in order to solve the rendering equation. A number of material models have been developed, since no model is currently able to reproduce the full range of available materials. Although only few material models have been widely adopted in current rendering systems, the lack of standardisation causes several issues in the 3D modelling workflow, leading to a heavy tool dependency of material appearance. In industry, final decisions about products are often based on a virtual prototype, a crucial step for the production pipeline, usually developed by a collaborations among several departments, which exchange data. Unfortunately, exchanged data often tends to differ from the original, when imported into a different application. As a result, delivering consistent visual results requires time, labour and computational cost. This thesis begins with an examination of the current state of the art in material appearance representation and capture, in order to identify a suitable strategy to tackle material appearance consistency. Automatic solutions to this problem are suggested in this work, accounting for the constraints of real-world scenarios, where the only available information is a reference rendering and the renderer used to obtain it, with no access to the implementation of the shaders. In particular, two image-based frameworks are proposed, working under these constraints. The first one, validated by means of perceptual studies, is aimed to the remapping of BRDF parameters and useful when the parameters used for the reference rendering are available. The second one provides consistent material appearance across different renderers, even when the parameters used for the reference are unknown. It allows the selection of an arbitrary reference rendering tool, and manipulates the output of other renderers in order to be consistent with the reference.

# Research Contribution

All publications, along with the relative contributions of the collaborating authors, that have resulted from the research presented in this thesis are listed in the following.

**BRDF Representation and Acquisition** *D. Guarnera, G.C. Guarnera, A. Ghosh, C. Denk, M. Glencross* [5]. The content of this paper constitutes the background research on material representation and acquisition, which defines the domain of the problem addressed in this thesis and presented in Chapters 2 and 3. The author began the research and wrote a consistent draft. Dr. Guarnera contributed to the research and in writing the paper. Dr. Ghosh provided important feedback and ideas on how to improve the quality and usefulness of the paper. Dr. Glencross helped in undertanding how to address such a huge body of research and helped in writing the paper.

Given the importance of virtual material capture and representation in computer graphics, two peer-reviewed courses have been presented at the top venue Conferences SIGGRAPH Asia 2016 and SIGGRAPH 2017, based on the content of Chapters 2 and 3: **Capturing and Representing BRDFs for Virtual Reality** *D. Guarnera, G. C. Guarnera, A. Ghosh, I. Hall, M. Glencross* [7] and **Material Capture and Representation for Virtual Reality** *G. C. Guarnera, A. Ghosh, I. Hall, M. Glencross, D. Guarnera* [8].

**Perceptually Validated Cross-Renderer Analytical BRDF Parameter Remapping** *D. Guarnera, G. C. Guarnera, M. Toscani, M. Glencross, B. Li, G. Y. Hardeberg, K. Gegenfurtner* [9]. The content of this paper, currently under submission and reported in Chapter 5, describes a perceptually validated automatic image

based approach for BRDF remapping. The author and Dr. Guarnera worked on the initial idea and designed the framework. Dr. Toscani carried out together with the author and Dr. Guarnera the psychophysical scaling experiments. Dr. Glencross contributed in designing some of the perceptual studies, conducted and analysed by the author. Dr. Li, Prof. Hardeber and Prof. Gegenfurtner supervised the project and provided feedback.

**Consistent, Tool Independent Virtual Material Appearance** *D. Guarnera, G.C. Guarnera, C. Denk, M. Glencross* [10]. The content of this paper, reported in Chapter 6, describes an automatic image based approach for BRDF remapping, in order to deliver material consistency. The author had the initial idea, implemented a large portion of the software and wrote most of the paper. Dr. Guarnera implemented portions of the software and helped refining the paper. Dr. Glencross provided input into the evaluation and during discussion. Ms. Denk provided an initial set of renderings for use in experiments.

# Chapter 1

# Introduction

In Computer Graphics, a range of high quality renderers are available either commercially or freely available to create realistic high quality imagery. The automotive industry uses a range of tools throughout the lifecycle of a car to visualise from design through to product. The lifespan of a vehicle can be in excess of 20 years. Rendering algorithms and techniques to represent the visual accuracy of car paints, coatings, interior materials develop at a much faster pace. For this reason, there is an interest in the automotive sector to be able to transfer the appearance of renderings created from certain packages to that created by other packages. This thesis addresses and solves a consistent material appearance transfer between different material models and rendering platforms. Our developed techniques evaluated through a user study to show efficacy of the novel method.

### 1.0.1 Problem definition

Digital reproduction of objects and real world complex scenes remains a challenge in the field of Computer Graphics, where the ultimate goal is to render images which are indistinguishable from pictures of their real world counterparts under the same lighting and viewing conditions. Photorealism has been the eternal quest of Computer Graphics since the earliest computers started to create pictures. Ivan Sutherland's Ultimate display [11] promised a level of realism indistinguishable from the real world. There are two parts to creating high quality renderings, the first part is acquisition of 3D geometry and the second part is acquiring re-

(a) Photo by Benjamin Child on Unsplash.com

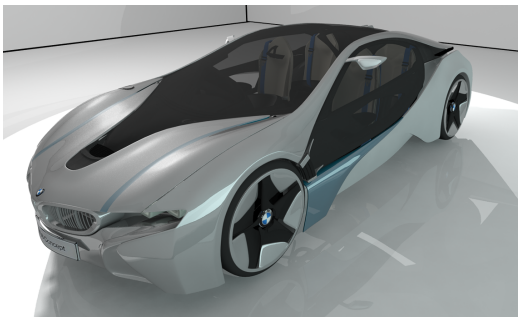

(b) Rendering in Autodesk 3ds Max by Dar'ya Guarnera

(c) Photo by Serge Kutuzov on Unsplash.com

Figure 1.1: A few examples of the range of images used in automotive industry, which spans from photo retouching to fully rendered images with different degrees of photorealism, enhancing different aspect of material appearance, like glossiness for instance.

alistic material representations. Depending on the accuracy of the geometry, the accuracy of the material representation and the features of the renderer, images can look very different while still being photorealistic. An interesting question this leads to is which renderer is a gold standard and most accurately re-creates the visual appearance of materials and coatings used in the design of a vehicle? Since paints and coatings have very different material reflectance characteristics to leather and cloth, different shaders are commonly employed to emphasise different visual effects depending on the application for the images. For example, a marketing brochure may employ more glossy looking visual appearance compared with reality, whereas a design/pre-visualisation application may require a more realistic appearance compared with reference materials available. Interestingly, this means that what the automotive industry considers to be photorealistic spans a gamut of different requirements based on application needs (figure 1.1).

Together with this application driven subjective photorealism, material appear-

ance is affected by the specific algorithms supported by the renderer used, some of these are proprietary IP and only available in in-house software packages. Material appearance also depends on the representation used in a specific tool (which also depends on supported algorithms). For example not all available renderers support subsurface scattering effects [12], or microfacet models [5], and this means that both the material model used and the visual output will differ from other bespoke renderers supporting new algorithms. Even when two renderers support the same material model, the available material properties to manipulate the appearance might differ (see Figure 1.2). Finally, also the rendering algorithm can influence the final appearance: in fact even within the same renderer the same material slightly differ when rendered two times, due to Monte Carlo sampling and other stochastic components.

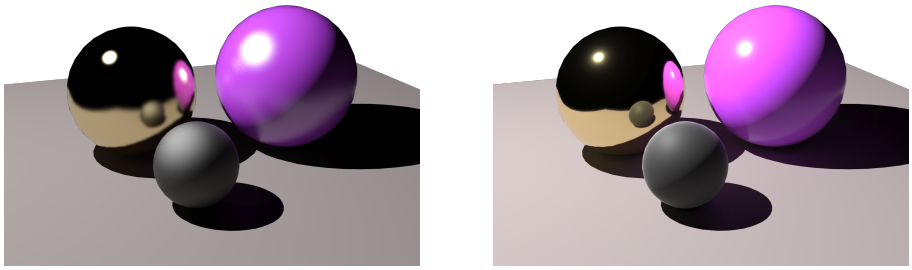

(a) (b)

Figure 1.2: The same scene, in which the same materials have been used, when rendered using two different renderers can look different due to the effect of the rendering algorithm. In (a) the NVIDIA® Mental Ray® plugin for 3ds Max® (Autodesk Inc.® ) is used, whereas in (b) the V-ray® plugin by Chaos Software® is used.

In Computer Graphics material appearance is still a challenge since:

- There is no widely adopted solution to represent a material;
- There is no standardised material model format
- There is little standardisation across renderers, with different renderers support subsets of material properties;
- Material models are hard to edit by artists.

As a consequence, material appearance is "tool dependent" and can require time, labour and computational cost to deliver consistent visual results. Most of the commercial 3D applications used in the automotive sector offer some level of artist directed control over material appearance. This enables artists to match the appearance of a synthetic representation of a material to a real world reference, but this is both labour intensive and subjective. Physically based renderers provide physically plausible material models, with the aim of eliminating the need to manually alter material appearance by artists, however, the resulting images do not always fit with the application needs for the image. Thus there is a competing demand in terms of the look and feel of the resulting rendering based on the target application of the image. The problem of unifying rendered appearance across renderers and across different material models are a key technical challenges addressed in this thesis.

### 1.0.2 Perceived Appearance

Our visual system is responsible for how we perceive what things are made of and the perception of material properties can involve all of our senses [13]. We extract material properties from the visual information we receive and we are good at recognizing and categorizing materials since our visual system can estimate the properties of materials from the retinal images. Numerous studies have shown that humans derive material properties through vision, which provides information on whether the material soft or hard, smooth or rough, real or fake, etc. Hence, human visual system perceives colour, lightness, roughness, transparency, etc. of a material, and also its feel. In order to recognise a material its overall colour distributions, intensities, contrasts or spatial attributes are not enough, in fact these are integrated by specific properties such as glossiness, surface roughness and translucency [14]. However, our knowledge of material perception is only in its infancy and there are many challenges are involved in material perception field [15, 16].

Visual perception may differ from subject to subject and individual differences arise for several reasons. Studies on differences in appearance perception show that we perceive things in a slightly different way from another observer due to individual differences in the colour perception and implicit assumptions about the

illumination [17]. Such differences may be unrelated across the different perceptual domains, thus complicating the analysis. Moreover, the human visual system is typically best at discriminating stimuli at the point it is adapted to, hence deviations from the present state is what we see best [18].

Given the importance of the context in which a scene is observed, which affects human perception, and the individual differences among subjects, it is easy to understand how the typical workflow of 3D modelling in industries such as automotive for instance, heavily based on the collaboration process between different digital artists belonging to different departments, is error prone. Departments such as design, visualisation and marketing have very different tasks, objectives and rendering scenarios. The broad range of modelling and rendering tools, each targeting different goals and requirements, provides the context in which artists adjust a rendering according to their own subjective visual perception, with the final result that renderings produced by artists in different departments might differ from each other, despite of the attempt to reuse digital assets as much as possible.

In this thesis the proposed techniques solving consistency of material appearance are validated by means of psychometric experiments using psycho toolbox and a common in Computer Graphics community method similarity judgement. Both experiments results show that out methods to unify material appearance are favorable among participants.

### 1.0.3 Transfer Methods

As previously stated, material modelling involves a great deal of manual effort at every stage, from the acquisition to the modelling and evening in post-rendering, since the lack of support to some feature might require some additional touch up, for instance with a photo editing software or with more sophisticated appearance transfer techniques.

Unfortunately, photo editing tools currently do not allow any editing of the material appearance at a BRDF level beyond recolouring or blurring, since modifications to the kind of reflection depicted in an image, or to the contribution of the Fresnel effect (which also determines the appearance of the silhouette of the objects), can only be done by manually painting pixel by pixel the appropriate

areas.

Appearance transfer and manipulation methods aim to provide the user with the ability to edit the appearance of materials, either by manipulating the pixels of the rendered image or photograph, by acting on the parameters of a given BRDF model before rendering an image, or by manipulating the appearance of the whole image using statistics derived from similar images.

In this thesis new approaches are proposed to characterise the statistics that represent the visual appearance of the output of a chosen renderer and map this to the visual appearance of another renderer. The benefit of these approaches in the automotive sector is that they reduce the dependence of images on the aforementioned differences.

### 1.0.4 Research Aims and Novel Contributions

Driven by the automotive industry, the focus of this project is on consistent materials representation. The following main aims and objectives can be identified:

- Eliminate the current tool-dependency of virtual materials, thus unifying material appearance across renderers and different virtual material representations and implementations.

- The solution needs to be automatic and must consider the constraints of real-world scenarios, where a material is exported from one software to an external renderer, or when a rendering created in a department making use of a specific tool is sent to a different one where such tool is not available.

In such scenarios there is only access to the final rendering (reference or *source*) and renderer used for it, with no information about the actual implementation of the shaders.

My contribution addresses two different variants of the same problem:

(a) information about the parameters used for the source material model and renderer is available;

(b) No information about the parameters used for the source material model and renderer .

To address the first case (a), we propose an image-based method for the remapping of BRDF parameters, based on a genetic algorithm optimisation and driven by a computational metric defined in the image space. Input to out method is a set of images of a scene, rendered with the source BRDF model, and the list of BRDF parameters used for them. The output of our method is a mapping from the source BRDF parameter space to the target BRDF model parameter space, such that a rendering with the remapped target parameters is visually undistinguishable from the rendering of the same scene with the source model. Our method, which works under the constraint of no knowledge of the model implementation, is described in Chapter 5. The robustness of the proposed method has been evaluated on a set of different material models, through both objective and subjective evaluation, the latter based on several user studies and psychometric scaling experiments. The diagram of the proposed solution is reported in Figure 1.3.

The scenario described in (b), in which the source parameters are not available, is addressed in Chapter 6. We propose a material appearance hallucination method aimed to deliver a tool-independent material appearance. Our method is based on a statistical characterisation of the output of different renderers, performed on a dataset tailored for the automotive industry (Figure 1.4). Such characterisation allows us to derive a set of correction functions, used to make output of the characterized renderers visually uniform to an arbitrarily chosen reference renderer. The application part of our framework takes in input an image rendered (or to be rendered) with any of the characterised renderers, and follows the diagram reported in Figure 1.5. The path taken depends on whether the input scene is already rendered or yet to render, and depends also on the information available about the materials used in the scene and the environment map. The output of our framework is a rendering which visually matches the appearance of the same scene as produced by the reference renderer.

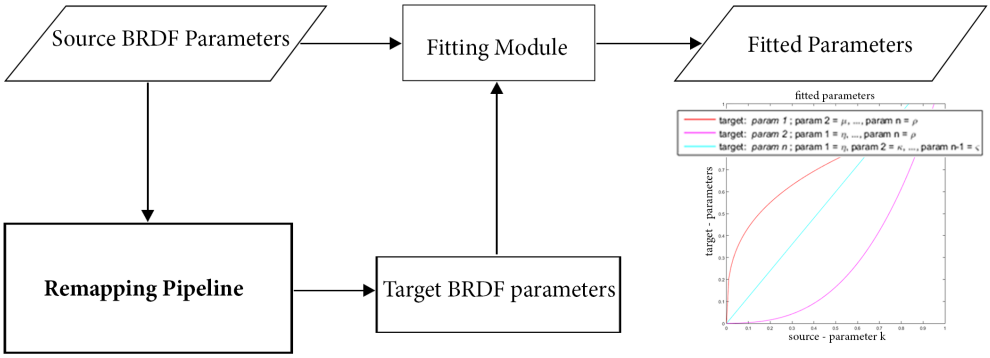


Figure 1.3: Diagram overview of our BRDF parameters remapping framework (Chapter 5).

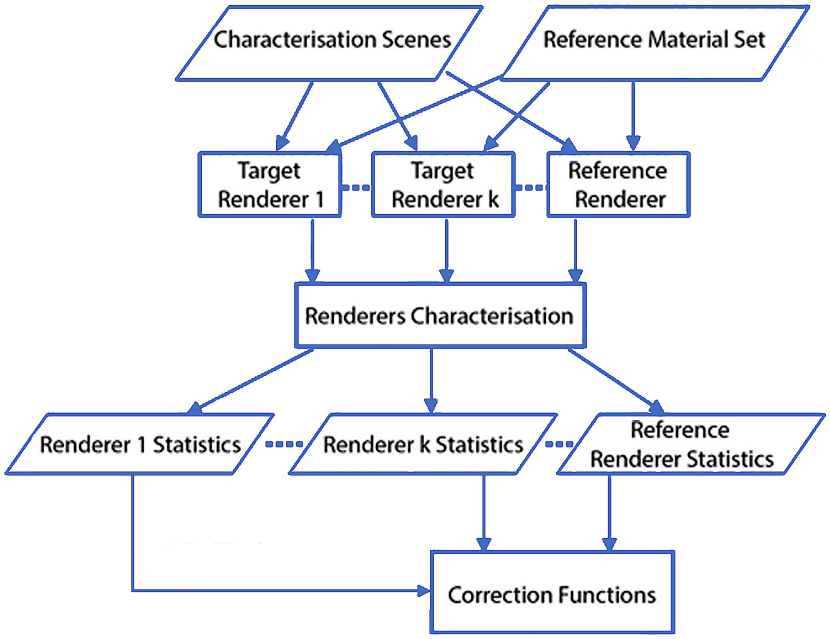


Figure 1.4: Diagram overview of renderers statistical analysis and characterisation part for tool independent material appearance framework (Chapter 6).

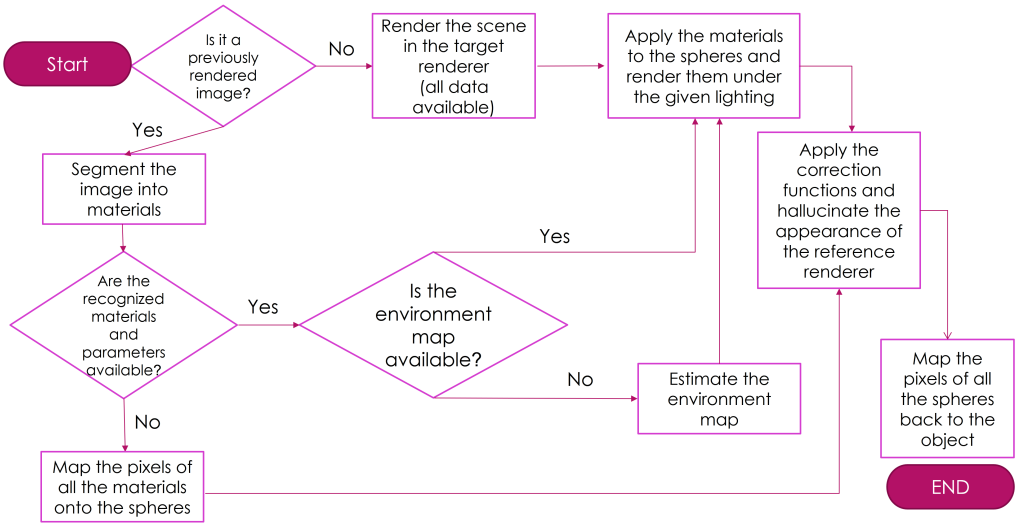


Figure 1.5: Diagram of the proposed framework application, taking in to account all possible scenarios, for consistent material appearance across different renderers (Chapter 6).

### 1.0.5 Thesis Organisation

This thesis begins with an overview of material appearance representation and acquisition methods, which define the domain of the problem addressed in this work. Existing techniques currently used to edit and manipulate the appearance of digital materials will also be described, before explaining in details the characteristics and achievements of this work.

Chapter 2 and Chapter 3 provide the reader with important background in material representation and acquisition in computer graphics and related fields. In particular Chapter 2 covers existing surface reflectance models in Computer Graphics, in order to explore the great variety of characteristic of the models currently used and understand their benefits and limitations; Chapter 3 describes the taxonomy of the available reflectance acquisition setups, defined in terms of the most important component, which also roughly defines benefits and limitations. Reflectance acquisition setups and reflectance models are closely related, since most reflectance models have been derived from acquired data and many acquisition setups make use of reflectance models to simplify the reflectance measurement task.

In Chapter 4 previous methods to manipulate the appearance of materials are described. They generally work either by manipulating the pixels of the rendered image or photograph, by acting on the parameters of a given BRDF model before rendering an image, or by manipulating the appearance of the whole image using statistics derived from similar images.

Chapter 5 details the proposed BRDF parameters remapping framework within a "minimum-ignorance" scenario, when at least the knowledge of the parameters used to create the reference rendering are available. A numerical validation of the remapping framework on dielectric and conductor materials is supported by a set of user studies.

Chapter 6 describes a solution for the scenario defined of "maximum-ignorance", when no information about the parameters is available. A general solution allowing to uniform the output of one renderer to an arbitrary selected "golden standard" is reported, based on a statistical characterisation of renderers on a benchmark dataset.

Finally, Chapter 7 concludes the thesis with a summary of all the material presented and discusses future directions of investigation.

# Chapter 2

# Bidirectional Reflectance Distribution Function (BRDF)

Material representation is an active research area, with more and more solutions being developed in order to represent materials accurately. A very common way to represent how an opaque, homogeneous surface interacts with the light is through the Bidirectional Reflectance-Distribution Function (BRDF) and a number of different BRDF models has been developed.

In this Chapter we report a survey of existing BRDF models in Computer Graphics. A careful exploration of BRDF models is pivotal in order to achieve the necessary background for the set of experiments described in Chapters 5 and 6, where the aim is to find analytical relations among different models and different renderers, by means of appearance matching methods. Such techniques heavily rely on prior knowledge about BRDF models, which we describe in this Chapter, in terms of parameters and properties accounted for; a table reported at the end of this Chapter (Table 2.1) serves as a short summary of the main characteristic of each BRDF model.

## 2.1 BRDF - Introduction and definitions

As discussed in the previous Chapter, one of the possible ways to represent the way an opaque, homogeneous material interacts with the light is through the

BRDF (Bidirectional Reflectance-Distribution Function), a radiometric function currently used to varying levels of accuracy in photorealistic rendering systems. It describes, in the general case, how incident energy redirected in all directions across a hemisphere above the surface. Historically, the BRDF was defined and suggested over the more generalised BSSRDF (Bidirectional Scattering- Surface Reflectance-Distribution function) [19] by Nicodemus [20], as a simplified reflectance representation for opaque surfaces: the BRDF assumes that light entering a material leaves the material at the same position, whereas the BSSRDF can describe light transport between any two incident rays on a surface. Many common translucent materials like milk, skin and alabaster cannot be represented by a BRDF since they are characterised by their subsurface scattering behaviour that smooth the surface details, with the light shining through them [21]. These materials are expensive to measure and render however many techniques have been proposed [19], [22], [23], [24], [25], [26].

The BRDF is defined as the ratio of the reflected radiance $L_r$ to incident irradiance (*i.e.* the incident flux per unit area of the surface):

$$f_r(\mathbf{v_i}, \mathbf{v_r}) = \frac{dL_r(\mathbf{v_r})}{dE_i(\mathbf{v_i})} = \frac{dL_r(\mathbf{v_r})}{L_i(\mathbf{v_i})\cos\theta_i d\omega_i} \tag{2.1.1}$$

where $\mathbf{v_i}$ and $\mathbf{v_r}$ are vectors describing the incident $(i)$ and exitant $(r)$ directions. By taking into account the incident radiance $L_i$ (*i.e.* the reflected flux per unit area per unit solid angle) instead of the $E_i$, thus considering the solid angle around the incident lighting direction and the cosine of the angle between the latter and the surface normal, we can write the previous Equation in a different form, which allows to understand how the units of a BRDF are inverse steradian $[1/sr]$:

$$f_r(\mathbf{v_i}, \mathbf{v_r}) = \frac{dL_r(\mathbf{v_r})}{L_i(\mathbf{v_i})\cos\theta_i d\omega_i} \tag{2.1.2}$$

.

Researchers have measured hundreds of BRDFs, suggested implementation techniques and allowed user input to edit and enhance materials. Recent implementations have expanded material libraries, but have not improved significantly upon material representation efficiency. However,the uptake of acquired models

has not been widespread across rendering packages due to their data and storage requirements.

To understand the way the BRDF is parameterised, let's take into consideration a point $p$ on a surface and the surface normal $\mathbf{n}$ at that specific location on the surface; on the plane tangent to the surface in $p$ we fix a reference direction $\mathbf{t}$, called tangent direction, and its perpendicular direction $\mathbf{b}$ on the plane: $\mathbf{n} \times \mathbf{t} \times \mathbf{b}$ define a local reference frame. Once we set the incoming light direction and the outgoing direction (viewing direction), the angle between the surface normal and the viewing direction is called $\theta_i$, similarly the angle between the surface normal and the outgoing direction is called $\theta_r$. If we take the projection of the viewing direction on the tangent plane, the angles between the tangent direction and the projection of the incoming direction are called respectively $\phi_i$, and $\phi_r$.

Figure 2.1(a) shows the geometry of the BRDF and the vectors used for parameterisations:

- $\mathbf{n}$ is the normal at a specific point $p$ on the surface
- $\mathbf{t}$ is the tangent vector. It is perpendicular to the normal $\mathbf{n}$ and hence it is tangent to the surface at $p$.
- $\mathbf{b}$ is the bi-tangent vector, defined as $\mathbf{b} = \mathbf{n} \times \mathbf{t}$. In literature it is also named as binormal vector.
- $\mathbf{h}$ is the halfway vector [27], defined as: $\mathbf{h} = \frac{(\mathbf{v_i}+\mathbf{v_r})}{\|\mathbf{v_i}+\mathbf{v_r}\|}$

Another very common way to parameterise the BRDF is the half way $\mathbf{h}$ vector shown in Figure 2.1(b), defined by the normalised vector sum of the incoming and outgoing directions. In this case we are taking into account the angle between the surface normal $\mathbf{n}$ and the halfway vector $\mathbf{h}$. This has important implications in the way the obtained data can be stored, compressed and can speed up computation of specific models. The use of the halfway vector enables another possibility to define a local reference frame, in which one of the axis is aligned with $\mathbf{h}$ and the other two are given by $\mathbf{b'} = \frac{(\mathbf{n}\times\mathbf{h})}{\|\mathbf{n}\times\mathbf{h}\|}$ and $\mathbf{t'} = \mathbf{b'} \times \mathbf{h}$.

There exist other coordinate system and parameterisation, especially suited for dimensionality reduction of some isotropic BRDF models, for instance the

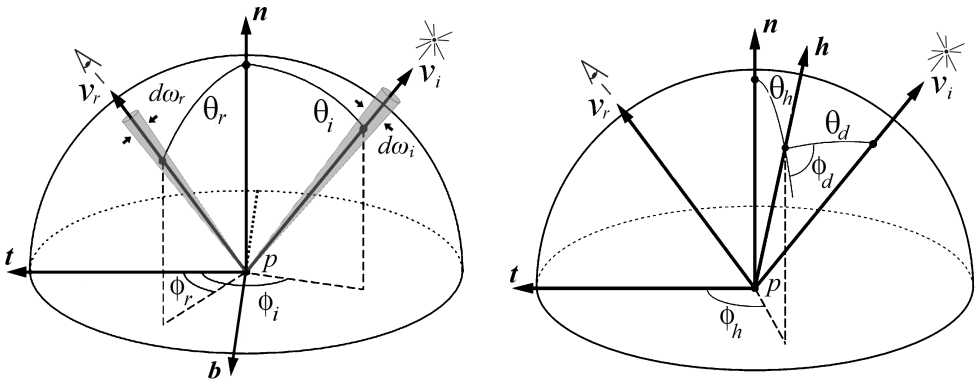


(a) Geometry of the BRDF (b) Halfway vector parameterisation

Figure 2.1: Geometry of the BRDF and halfway vector parameterisation.

barycentric coordinate system with respect to a triangular support proposed by Stark *et al.* [28], or the hybrid model described by Barla *et al.* which could lead to a better repartition of samples to cover most of the effects of materials [29].

## 2.2 Properties of the BRDF

There are many reflectance models that are simplified subsets of the BRDF function. One of the simplest reflectance models is the Lambertian model, which represents the perfect diffuse reflectance and is often used in many interactive applications, since it requires no recalculation with the change of viewing direction. The model simply assumes that the surface reflects light uniformly in all directions with the same radiance (see Figure 2.2, in yellow), constant with $v_r$ unlike other BRDF models: $f_r(\mathbf{v_i}, \mathbf{v_r}) = \rho_d/\pi$, where $\rho_d$ is the diffuse albedo.

In the case of a pure specular BRDF all the light is reflected in a single direction, for a given incident direction (see Figure 2.2, in light blue). In fact, light that is incident within a differential solid angle $d\omega_i$ from direction $(\theta_i, \phi_i)$ is reflected in a differential solid angle $\omega_r$ in direction $(\theta_i, \phi_i + \pi)$, hence the pure specular BRDF can be formalised with a double Dirac delta function: $f_r(\mathbf{v_i}, \mathbf{v_r}) = \rho_s \delta(\theta_i - \theta_r)\,\delta(\phi_i + \pi - \phi_r)$, where $\rho_s = L_r/L_i$ is the specular albedo. Perfect specularity is valid only for highly polished mirrors and metals.

Surfaces not perfectly smooth, which have some roughness at the micro-geometry

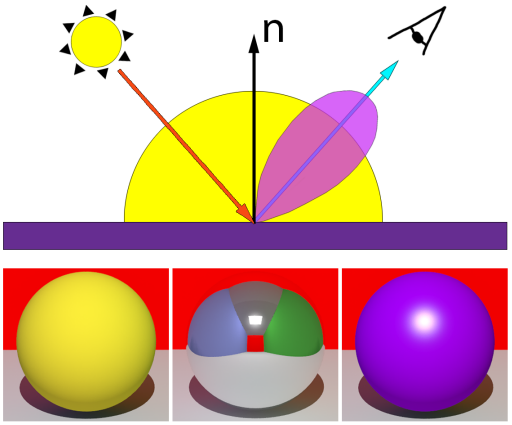


Figure 2.2: Basic reflectance models of the incoming light (in orange): perfect diffuse (yellow), glossy (purple) and perfect specular (light blue). Renderings of diffuse, glossy and specular spheres are shown, placed inside a Cornell box [1].

level, have a glossy appearance and show broader highlights, other than specular reflections (see Figure 2.2).

Some materials, like the surface of the moon or some biological tissues, show a phenomenon called retro-reflection in which light is scattered not only in the forward direction but also in the direction of the illuminant.

A BRDF should respect some basic physical properties, namely non-negativity, reciprocity and energy conservation:

- non-negativity: the BRDF is a non-negative function, hence for any pair of incident and outgoing direction $f_r(\mathbf{v_r}, \mathbf{v_i}) \geq 0$;
- the Helmholtz reciprocity principle states that the light path is reversible, for any pair of incident and outgoing direction: $f_r(\mathbf{v_i}, \mathbf{v_r}) = f_r(\mathbf{v_r}, \mathbf{v_i})$. This principle holds only for corresponding states of polarisation for incident and emerging fluxes, whereas large discrepancies might occur for non-corresponding states of polarisation [30]. In designing a rendering system possible non-reciprocity should be taken into account [31].
- Energy conservation assumes that the energy reflected cannot exceed incident energy [32]: $L_r \leq E_i$ hence over the unit hemisphere $\Omega_+$ above the

surface

$$\forall \mathbf{v_i}, \int_{\Omega_+} f_r(\mathbf{v_i}, \mathbf{v_r})(\mathbf{v_r} \cdot \mathbf{n}) d\omega_r \leq 1 \tag{2.2.1}$$

BRDFs can be classified by taking into account the characteristics of the material to represent:

- *Isotropic* BRDFs are able to represent materials whose reflection does not depend on the orientation of the surface, since the reflectance properties are invariant to rotations of the surface around **n**.
- *Anisotropic* BRDFs can describe materials whose reflection change with respect to rotation of the surface around **n**; this class includes materials like brushed metal, satin, velvet and hair.

The Fresnel effect predicts the fraction of power which is reflected and transmitted and has a great impact on the appearance. Many basic BRDF models have lost importance in the context of physically based modelling because they do not account for a Fresnel term. For conductive materials, like metals, the fraction of light reflected by pure specular reflection is roughly constant for all angles of incidence, whereas for non-conductive materials (dielectrics), the amount of light reflected increases at grazing angles; see (Figure 2.3) for a comparative example of the behaviour of metals and dielectrics. The fraction of light reflected is called Fresnel reflectance, which can be obtained from the solution of Maxwell's equations and depends also on the polarisation state of the incident light. For unpolarised light, the Fresnel reflectance $\Im$ at the interface between the surface and the air is given by

$$\Im(\eta, \theta_i, \theta_t) = \frac{1}{2}\left[\left(\frac{\eta\cos\theta_i - \cos\theta_t}{\eta\cos\theta_i + \cos\theta_t}\right)^2 + \left(\frac{\cos\theta_i - \eta\cos\theta_t}{\cos\theta_i + \cos\eta\theta_t}\right)^2\right], \tag{2.2.2}$$

where $\eta$ is the index of refraction of the surface and $\theta_t$ is the angle of transmission. In Computer Graphics, it is very common to use Schlick's approximation of the Fresnel reflectance [33]: $\Im(\theta) = \Im(0) + (1 - \Im(0))(1 - \cos(\theta))^5$; in Section 2.3 we will generally use the symbol $\Im$ to refer either to the exact Fresnel reflectance or one of its approximations.

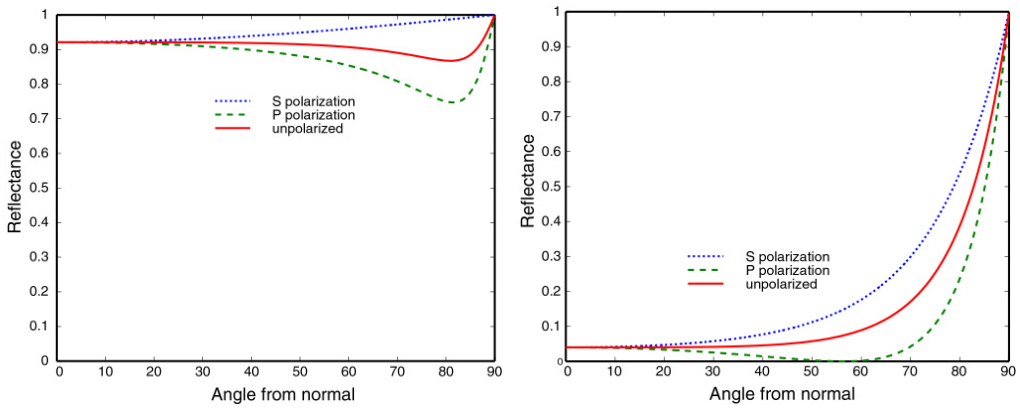


(a) Metal (aluminium) (b) Dielectric (plastic)

Figure 2.3: Fresnel reflectance for metals (a) and dielectrics (b).

## 2.3 BRDF Models

In the last 40 years many material models are been proposed and some of BRDFs are able to describe a wide subset of a material properties we have mentioned. Generally they can be classified into 3 families (Figure 2.4) which have different purposes and different way of calculating interaction with light:

- Phenomenological
- Physically based
- Data driven

Many medium, such as hair, fur, cloth and knitwear are difficult to describe by a surface model. These materials, and objects with highly complex boundary, are better described by volumetric appearance models [34–36], in particular for closer viewing distance, whereas BRDFs can be used from farther away. Jakob *et al.* [37] introduced a generalisation to anisotropic scattering structures, exploited also for volumes acquired by CT scans [38]. More recently, collections of individual fibers have been used for fabric representation [39]. In this thesis I focus on representation and acquisition of surface reflectance, hence I will not further discuss volumetric representation.

An very important aspect is the practicality of a model in a rendering system, which requires a suitable technique for importance sampling. When calculating

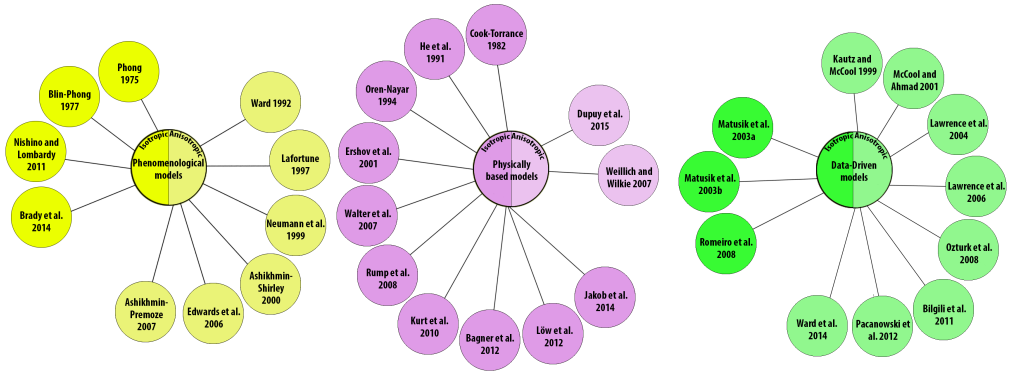


Figure 2.4: Taxonomy of Reflectance models colour coded for isotropic and anisotropic representations

the radiance direction of a surface in a scene, accounting for the contribution of light from all possible directions is expensive to compute, therefore Monte Carlo techniques are used to estimate the values with fewer samples [40], based on a stochastic process. However, the number of samples should be sufficient to produce consistent estimations, otherwise the results will vary significantly. Importance sampling can be used to reduce sample variance [41], by distributing samples according to the known elements, either taking into account the reflection model in use or the incident light [42].

All material models have been developed for general material representation but those models still require further development as the parameters are not as intuitive and controllable as they are meant to be. In the Figure 2.4 the taxonomy of reflectance models is reported.

## 2.4 Phenomenological models

Phenomenological models are entirely based on reflectance data, which is fitted to analytical formulas, thus approximating the reflectance and reproducing characteristics of real world materials, but they will not necessarily appear realistic unless placed in an accurately simulated environment for such a model.

### 2.4.1 Phenomenological models for Isotropic materials

The simplest BRDF model is the Phong [43]. The model is useful for isotropic materials with slightly rough surface. The specular reflection computed using specular constant and the dot product between the reflected direction and the mirrored reflection of the incoming light. You take the incoming direction of the light you compute the mirrored direction of the incoming direction and then you compute the dot product between the reflected incoming direction and the outgoing direction. The cosine of this angle razed to some exponent n will give you the final reflection and it controls the width of the specular highlight, higher values of n gives more specular appearance. and useful for isotropic materials $f_r(\mathbf{v_i}, \mathbf{v_r}) = k_s(\mathbf{v_r} \cdot \mathbf{r_{vi}})^n$, where $k_s$ is a specular constant in the range $[0, \infty]$, $\mathbf{r_{vi}}$ is the direction of $\mathbf{v_i}$ after being perfectly reflected and $n$ controls the shape of the specular highlight; in Figure 2.5 we report some renderings using the Phong model, generated by the Mitsuba renderer [44], for different values on the exponent $n$, for both point light illumination and environmental map illumination. This model does not take into account energy conservation nor reciprocity. Moreover, it does not capture the reflection behaviour of real surfaces at grazing angles. It is not normalised, however some normalisation factors for cosine lobes have been proposed, either based on double-axis moments [45] or with the simpler option of a power series in $(\mathbf{n} \cdot \mathbf{h})$ with a suitable sequence of exponents [46].

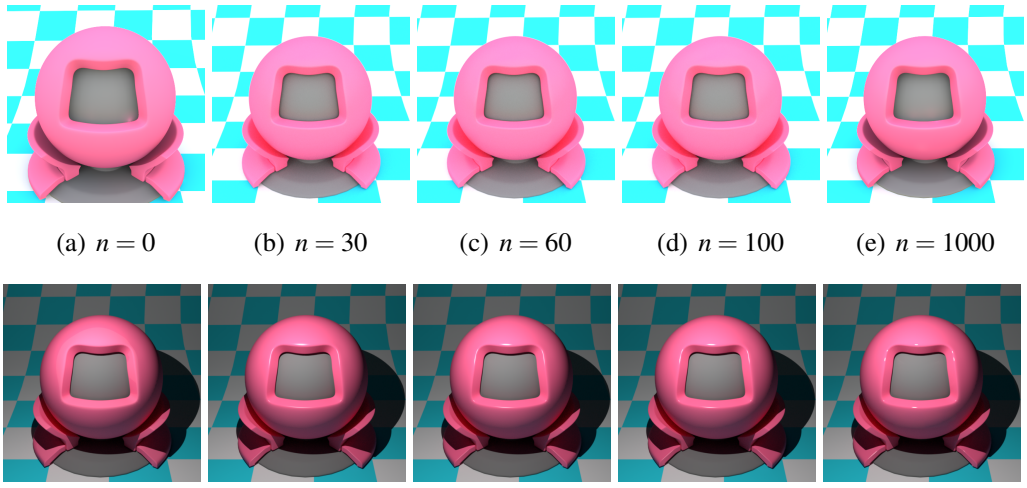

(a) $n = 0$ (b) $n = 30$ (c) $n = 60$ (d) $n = 100$ (e) $n = 1000$

(f) $n = 0$ (g) $n = 30$ (h) $n = 60$ (i) $n = 100$ (j) $n = 1000$

Figure 2.5: Phong model. The top row reports a sequence of renderings for increasing values of $n$, under environment map lighting; from left to right: $n = 30, 60, 100, 1000$. The bottom row shows the same sequence of exponents, with a point light illumination.

Later Blin-Phong [47] suggested a correction to their model in which it is still the cosine lobe model but in this case they calculated cosine, the dot product of the surface normal and the halfway vector and this is due to the fact that in those years it was costly to calculate the mirror direction of the incoming light so in this way the computation is faster. Although it has been used as the default shading model for OpenGL and Direct3D until recent times, it shares the same limitations of the Phong model, being physically not plausible does not fulfil reciprocity, no energy conservation, no Fresnel effect and not able to capture metallic and mirrored appearance.

$$f_r(\mathbf{v_i}, \mathbf{v_r}) = k_s(\mathbf{n} \cdot \mathbf{h})^n \tag{2.4.1}$$

Since it follows the cosine function, if $n$ goes to infinity the reflected radiance and the albedo converges to zero towards grazing angles.

Nishino and Lombardi suggested model based on the directional statistics in particular the reflection of the surface described by using a multiple lobes which is automatically computed by expectation maximization algorithm so this algorithm tells you how many lobes to use. Each lobe is a statistical density function distribution which means given an input incoming direction the model will give you an output a distribution of possible reflected directions which shape

the reflected lobe, by putting them together we obtain lobes with a more complex function with parameters which controls the shapes which are the height of the lobe and the width of the lobe. A statistical distribution on a unit hemisphere, *i.e.* a probability density function that takes in input a direction for $\mathbf{v_i}$ and returns a distribution of directions for $\mathbf{v_i}$, called Hemispherical Exponential Power Distribution (Hemi-EPD). The Hemi-EPD constitutes a basis for the entire BRDF, which can be modeled as mixtures of Hemi-EPDs, one for each of its 2D slices. The BRDF is parameterised using the halfway vector and difference angle, by dropping the dependence on $\phi_h$ because of the isotropy. The expression for the Hemi-EPDs, given the parameters $\Theta = \kappa, \gamma$ and the normalization factor $C$, is $p(\theta_h|\theta_d, \Theta) = C(\Theta)(e^{\kappa \cos^{\gamma} \theta_h} - 1)$. The optimal number $k$ of lobes is automatically determined by an Expectation-Maximisation algorithm, which tests different numbers of lobes with a statistical measure; as long as the condition $\sum_1^k 1/C(\Theta_k) \leq 1$ holds, energy conservation is guaranteed.

Brady et al 2014 [48] suggested a machine learning approach to learn new analytical models to represent BRDFs he used a training set which is a subset of the MERL dataset using a few seeds which are the starting point of the genetic programming this few seeds are very simple BRDF models. A few basic BRDF models are used as a starting point (seeds), on which symbolic transformations are applied. The fitness function calculates the residual error of each variant after fitting the free parameters to the training set of isotropic materials, a subset of the MERL-MIT database [4]. The random search is heuristic-based, trying to adapt the starting models to the measured ground truth data and tends to produce a large set of candidates expressions. To allow a better exploration of the search space some suboptimal variations that increase the error are allowed and for the same purpose an island model genetic algorithm is used, allowing only sporadic interactions between sub-populations. The grammar does not guarantee that the resulting models respect energy conservation and reciprocity, hence these properties need to be taken into account by the fitness function; in [48] is reported a table with some variants for which the properties have been numerically verified.

### 2.4.2 Phenomenological models for Anisotropic materials

The Ward reflectance model [49] is able to represent both isotropic and anisotropic reflection; it combines specular and diffuse components of reflectance, representing specular peaks through Gaussian distributions. It was specifically designed to easily fit measured BRDFs, which have been used for validation. The model specifies an efficient for Monte Carlo sampling. The Ward model has four parameters, which can be set independently, therefore it can be fitted to a large class of measured data. The anisotropic model makes use of the two parameters $\alpha_x$ and $\alpha_y$ to control the width of the gaussian lobe in the two principal directions of anisotropy:

$$f_r(\mathbf{v_i}, \mathbf{v_r}) = \frac{\rho_d}{\pi} + \frac{\rho_s}{\sqrt{\cos(\theta_i)\cos(\theta_r)}} \cdot \frac{e^{-\tan^2(\theta_h)\left(\frac{\cos^2\theta_h}{{\alpha_x}^2} + \frac{\sin^2\theta_h}{{\alpha_y}^2}\right)}}{4\pi\alpha_x\alpha_y} \tag{2.4.2}$$

where $\rho_s$ controls the magnitude of the lobe and $4\pi\alpha^2$ is a normalisation factor. The isotropic Ward model is obtained by setting $\alpha_x = \alpha_y$.

The model does not obey the principle of energy conservation at grazing angles, which has been investigated in [50], [51], [2]. A different normalisation factor has been proposed in [51] to prevent numerical instabilities and to correct the loss of energy at flat angles, specifically $(4\cos(\theta_i)\cos(\theta_r))$ instead of $\left(4\sqrt{\cos(\theta_i)\cos(\theta_r)}\right)$, however it shares the problem of diverging to infinity with the original Ward model. A new physically plausible version of the model has been proposed in [2], which meets the energy conservation principle even at grazing angles by using the following normalisation factor:

$$\frac{2\left(1 + \cos\theta_i\cos\theta_r + \sin\theta_i\sin\theta_r\cos\phi_r - \phi_r\right)}{\left(\cos\theta_i\cos\theta_r\right)^4}. \tag{2.4.3}$$

In Figure 2.6 we report a set of renderings of the Ward energy conserving variant described in [2], using the Mitsuba renderer [44], for different values of specular reflectance, for both isotropic and anisotropic materials.

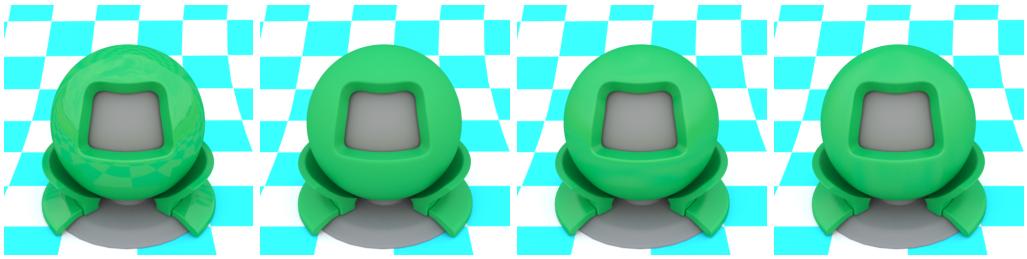

(a) s = 0.2, r = 0.01 (b) s = 0.2, r = 0.5 (c) s = 0.2, r = {0.3, 0.05} (d) s = 0.2, r = {0.05, 0.3}

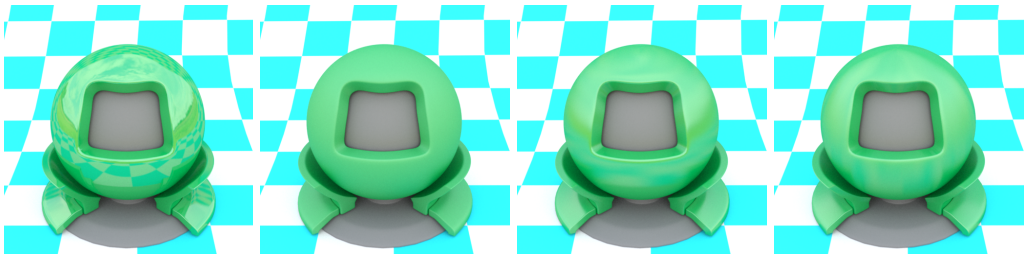

(e) s = 0.4, r = 0.01 (f) s = 0.4, r = 0.01 (g) s = 0.2, r = {0.3, 0.05} (h) s = 0.2, r = {0.05, 0.3}

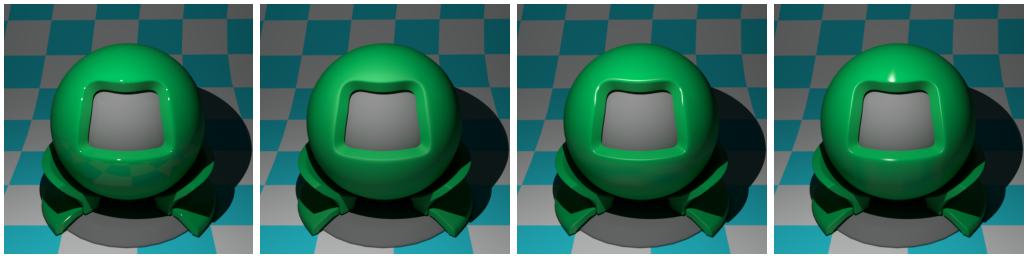

(i) s = 0.2, r = 0.01 (j) s = 0.2, r = 0.5 (k) s = 0.2, r = {0.3, 0.05} (l) s = 0.2, r = {0.05, 0.3}

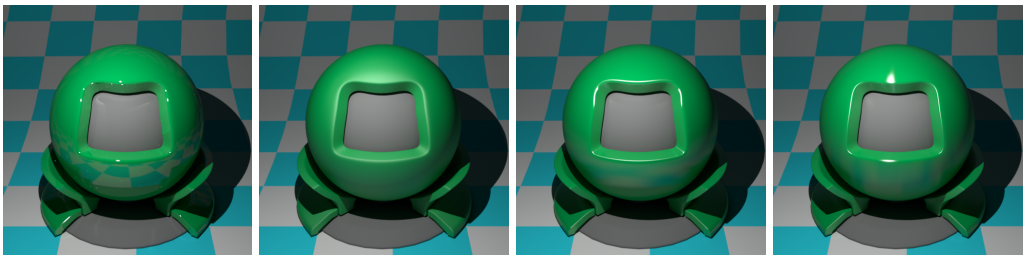

(m) s = 0.4, r = 0.01 (n) s = 0.4, r = 0.01 (o) s = 0.2, r = {0.3, 0.05} (p) s = 0.2, r = {0.05, 0.3}

Figure 2.6: Energy conserving Ward model variant [2]. The first two columns refer to an isotropic material, whereas the last two columns refer to an anisotropic material. The top two rows reports a sequence of renderings under environment map lighting. The bottom two rows show the same sequence under a point light illumination. The parameters used are reported in brackets: the first value (s) refers to the specular reflectance, the second value (r) to the isotropic roughness, or in case of anisotropic material the second value refers to the roughness in the tangent direction, the second one to the bitangent direction.

The Lafortune [52] model is a flexible, empirical model designed to fit measurements from real surfaces and compactly represent them [53]. The model is a generalisation of the cosine lobe model with multiple steerable lobes, based on the Phong model. The primitive functions obey the Energy Conservation and Reciprocity principles. This model allows lobe specification on the surface in terms of shape and direction, by simply setting up to 3 parameters and an exponent:

$$f_r(\mathbf{v_i}, \mathbf{v_r}) = \frac{\rho_d}{\pi} + \sum_{l=1}^{N} \left(C_{x,l} v_{ix} v_{rx} + C_{y,l} v_{iy} v_{ry} + C_{z,l} v_{iz} v_{rz}\right)^{n_l} \tag{2.4.4}$$

where $N$ is the number of lobes, $C_x$, $C_y$, $C_z$ are parameters which absorb the specular albedo and control retro-reflections (by setting $C_x$, $C_y$ and $C_z$ to positive values), anisotropy (with $C_x \neq C_y$) and off-specular peaks (if $C_z$ is smaller than $-C_x = -C_y$). Lafortune′s reflection model can represent generalised diffuse reflectance as the model is able to reflect radiance evenly in all directions, by setting $C_x = C_y = 0$; the Lambertian model can be obtained by setting $N = 0$. A comparative study shows that the Lafortune model performs better than the Phong, Ward and He *et al.* models in representing measured BRDFs like white paper, rough plastic, rough aluminium and metal, since it was designed to fit almost any BRDF data [53].

Neumann *et al.* [50] proposed some modifications and correction factors for the reciprocal Phong [43], [54], Blinn [47] and Ward [49] models. The correction factors can be seen as shadowing and masking terms to make the models physically plausible. Moreover the modified models can be used to render metals and other specular objects and for each of them an importance sampling procedure is described.

The Ashikhmin-Shirley model is based on the Phong reflectance model [55]. The reflectance of the model changes with the view-point, hence at grazing angles the reflectance is specular and at normal angles the reflectance is diffuse. This model assumes micro facets with various angles and sizes [56] by generalising the types of microfacets and allows the expression of arbitrary angles.

The specular component of the BRDF is expressed by:

$$f_{r,s}(\mathbf{v_i}, \mathbf{v_r}) = \frac{\Im(\mathbf{v_i} \cdot \mathbf{h}) D(\mathbf{h})}{2(\mathbf{h} \cdot \mathbf{v_r}) \max(\mathbf{n} \cdot \mathbf{v_r}, \mathbf{n} \cdot \mathbf{v_i})} \tag{2.4.5}$$

where $\Im(\mathbf{v_i} \cdot \mathbf{h})$ is the Schlick′s approximation of the Fresnel term [33]. $D(\mathbf{h})$ is the distribution function of the microfacets, controlled by the parameters $e_x$ and $e_y$, the axes of an ellipse which orientates the halfway vector $\mathbf{h}$ of the microfacets respectively along the $X$ and $Y$ and thus defining the anisotropy:

$$D(\mathbf{h}) = \frac{\sqrt{(e_x+1)(e_y+1)}\,(\mathbf{h} \cdot \mathbf{n})^{e_x \cos^2(\phi_h) + e_y \sin^2(\phi_h)}}{4\pi}. \tag{2.4.6}$$

In order to preserve energy conservation and to model the behaviour of the surface's diffuse colour near the grazing angle, which disappears due to the increase in specular reflectance, instead of a Lambertian diffuse term an angle-dependent form of the diffuse component is reported. The expression is based on the consideration that the amount of energy for diffuse scattering is dependent on the total reflectance of the specular term at the incident angle:

$$f_{r,d}(\mathbf{v_i}, \mathbf{v_r}) = (1 - \rho_s)\, g(\mathbf{v_i}, \mathbf{v_r}, \mathbf{n})(28\rho_d)/(23\pi) \tag{2.4.7}$$

where $g(\mathbf{v_i}, \mathbf{v_r}, \mathbf{n}) = [1 - (1 - (\mathbf{n} \cdot \mathbf{v_i})/2)^5][1 - (1 - (\mathbf{n} \cdot \mathbf{v_r})/2)^5]$. The model is able to describe anisotropic reflections of two layered materials, such as varnished wood for example and it is physically plausible. A sampling method for Monte Carlo rendering is also provided, based on $D(\mathbf{h})$: it gives the probability density function $p(\mathbf{v_r}) = D(\mathbf{h})/4(\mathbf{v_i} \cdot \mathbf{h})$. In Figure 2.7 we report some renderings with the Ashikhmin-Shirley model for a conductor material (gold), showing the effect of different values of roughness, both in the isotropic and anisotropic case.

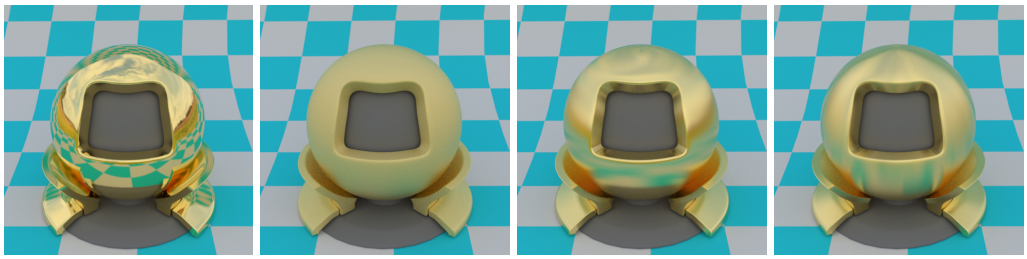

(a) roughness = 0.01 (b) roughness = 0.5 (c) roughness = {0.3, 0.05} (d) roughness = {0.05, 0.3}

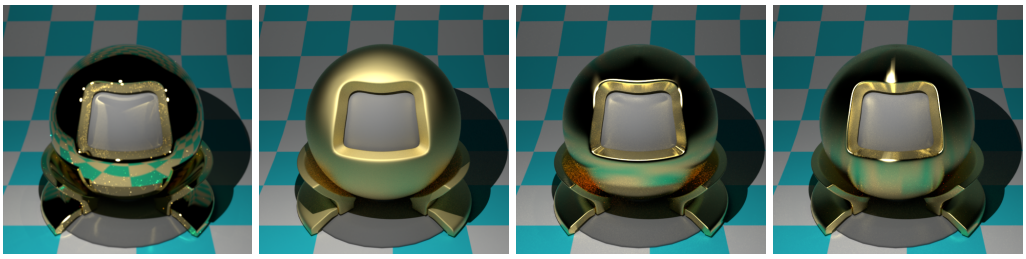

(e) roughness = 0.01 (f) roughness = 0.01 (g) roughness = {0.3, 0.05} (h) roughness = {0.05, 0.3}

Figure 2.7: Ashikhmin-Shirley BRDF model. The first two columns refer to an isotropic material, whereas the last two columns refer to an anisotropic material. The top rows reports a sequence of renderings under environment map lighting. The bottom two rows show the same sequence under a point light illumination. The parameters used are reported in brackets; in case of anisotropic material the values respectively refer to the roughness in the tangent direction and in the bitangent direction.

Edwards *et al.* proposed a framework for transforming the halfway vector $\mathbf{h}$ into different domains to enforce energy conservation but compromising reciprocity [57]. By writing Equation 2.2.1 in terms of $\forall \mathbf{v_r}$ and assuming that it satisfies an equality instead of an inequality, the function

$$Q(\mathbf{v_i}) = f_r(\mathbf{v_i}, \mathbf{v_r})(\mathbf{v_i} \cdot \mathbf{n}) \tag{2.4.8}$$

can be seen as a probability density function (PDF) over the set of incident directions $\mathbf{v_i}$ on the hemisphere $\Omega_+$. Since the PDF $Q(\mathbf{v_i})$ is related to a PDF $q(\mathbf{h})$ over halfway vectors by the formula $Q(\mathbf{v_i}) = q(\mathbf{h})/(4\mathbf{v_i} \cdot \mathbf{h})$, from Equation 2.4.8 the following expression for $f_r(\mathbf{v_i}, \mathbf{v_r})$, which conserves energy, can be derived:

$$f_r(\mathbf{v_i}, \mathbf{v_r}) = [q(\mathbf{h})] \,/\, [4(\mathbf{v_i} \cdot \mathbf{n})(\mathbf{v_i} \cdot \mathbf{h})]\,. \tag{2.4.9}$$

In the space of incident directions $\mathbf{v_i}$ it is difficult to formulate a PDF to describe

off-specular reflection and other phenomena. As for the halfway vector domain, near to grazing angles the set of allowable halfway vectors changes in a complicated way. If a new domain $D_h$ is defined, together with a PDF $p(\mathbf{l})$ and a bijection $f(\mathbf{h}) = \mathbf{l}$ between the set of halfway vectors $\mathbf{h} \in \Omega_+$ and the set of points $\mathbf{l} \in D_h$, by equating the differential probabilities between $D_h$ and $\Omega_+$ the following can be derived from Equation 2.4.9:

$$f_r(\mathbf{v_i}, \mathbf{v_r}) = [p(\mathbf{l})d\mu(\mathbf{l})] \,/\, [(4\mathbf{v_i} \cdot \mathbf{h})d\omega_h] \tag{2.4.10}$$

where $d\mu$ is the differential measure over $D_h$ and $p(\mathbf{l})d\mu(\mathbf{l}) = q(\mathbf{h})d\omega_h$. With this framework, a new domain can be defined given $\mathbf{v_r}$, by translating $\Omega_+$ so that the center of its base lies at the tip of $\mathbf{v_r}$. In this way, every point on the translated hemisphere corresponds to a not normalised halfway vector $\mathbf{h_u} = \mathbf{v_i} + \mathbf{v_r}$. The final step is the transformation of the vectors $\mathbf{h_u}$ to points $\mathbf{l}$ on the base of the hemisphere; if the local orientation of the surface is given by the orthogonal vectors $\mathbf{u}$ and $\mathbf{v}$, a point on the disk can be defined by the $(u, v)$ coordinates, hence the PDF $p(\mathbf{l})$ is two dimensional. A possibility is to scale the halfway vector until its tip lies in the base of the hemisphere, and the resulting energy conserving BRDF is:

$$f_r(\mathbf{v_i}, \mathbf{v_r}) = \frac{p(\mathbf{l})(\mathbf{v_r} \cdot \mathbf{n})^2}{4(\mathbf{v_i} \cdot \mathbf{n})(\mathbf{v_i} \cdot \mathbf{h})(\mathbf{h} \cdot \mathbf{n})^3} \tag{2.4.11}$$

since $\mathbf{l} = \frac{(\mathbf{v_r} \cdot \mathbf{n})}{(\mathbf{h_u} \cdot \mathbf{n})}\mathbf{h_u}$. This transform allows to describe retro-reflective materials by defining a PDF with high values near the center of the disk, which corresponds to a halfway vector in the retro-reflective direction; to specify a shiny BRDFs it is enough to define a PDF with high values near the origin of the $(u, v)$ space, which corresponds to $\mathbf{n}$ and gives pure specular reflection. To importance sample the BRDF to obtain $\mathbf{h}$ it is enough to generate a point $\mathbf{l}$ on the disk according to $p(\mathbf{l})$ and normalise. Alternatively, the orthogonal projection maps $\mathbf{h_u}$ to the disk along the direction of $\mathbf{n}$: $\mathbf{l} = \mathbf{h_u} - (\mathbf{v_i} \cdot \mathbf{n})\mathbf{n}$; it leads to a BRDF with narrower lobes, centered on the direction of perfect reflection. The resulting BRDF, suitable for data fitting, is given by:

$$f_r(\mathbf{v_i}, \mathbf{v_r}) = [1/(4(\mathbf{v_i} \cdot \mathbf{h})^2)]p(\mathbf{l}) \parallel \mathbf{v_i} + \mathbf{v_r} \parallel^2 \tag{2.4.12}$$

To importance sample the BRDF, once a sample $\mathbf{l}$ is generated according to $p(\mathbf{l})$, the unnormalised halfway vector is obtained from the expression $\mathbf{h_u} = \mathbf{l} + (\mathbf{v_i} \cdot \mathbf{n})\mathbf{n}$. Within the same framework two additional BRDF models are described: an empirical, energy-preserving BRDF with limited number of parameters and a BRDF model useful for data fitting, which does not preserve energy.

The Ashikhmin-Premoze model [58], or d-BRDF, follows [59] and [60] microfacet theory and it is based on the earlier Ashikhmin-Shirley model [55], with a simplified process of fitting BRDF models to measured data; an efficient sampling technique is also suggested. The Ashikhmin-Premoze model combines an analytic model with a data driven distribution and also discusses how to fit backscattering measurements to the model [61]. The model allows the use of an arbitrary normalised function $p(\mathbf{h})$, hence specular highlights can be easily adjusted since their shape depends directly on the distribution. The *max* term in Equation 2.4.5, which causes colour banding artifacts as observed in [58], is replaced with a smoother term $(\mathbf{v_i} \cdot \mathbf{n}) + (\mathbf{v_r} \cdot \mathbf{n}) - (\mathbf{v_i} \cdot \mathbf{n})(\mathbf{v_r} \cdot \mathbf{n})$. An additional modification is to exclude the $(\mathbf{h} \cdot \mathbf{v_r})$ term to improve the appearance matching with real world materials. The resulting expression for the specular term, which is reciprocal and non negative for any non-negative $p(\mathbf{h})$, can be written as:

$$f_{r,s}(\mathbf{v_i}, \mathbf{v_r}) = \frac{\Im(\mathbf{v_i} \cdot \mathbf{h})\, p(\mathbf{h}) k_s}{(\mathbf{v_i} \cdot \mathbf{n}) + (\mathbf{v_r} \cdot \mathbf{n}) - (\mathbf{v_i} \cdot \mathbf{n})(\mathbf{v_r} \cdot \mathbf{n})} \tag{2.4.13}$$

where $k_s$ is a scaling constant which needs to be chosen in order to fulfill energy conservation. The d-BRDF model improves representation of material reflectance at grazing angles and enables more realistic material appearance, however some effects like retro-reflection cannot be modeled properly. The values of the Fresnel parameters can lie outside the actual range for a given materia, hence they do not have a physical meaning.

## 2.5 Physically based models

Physically based models are based on Physics and Optics, with the assumption that the surface is rough at a fine scale, therefore described by a collection of micro facets with some distribution $D$ of size and direction. Usually they are represented

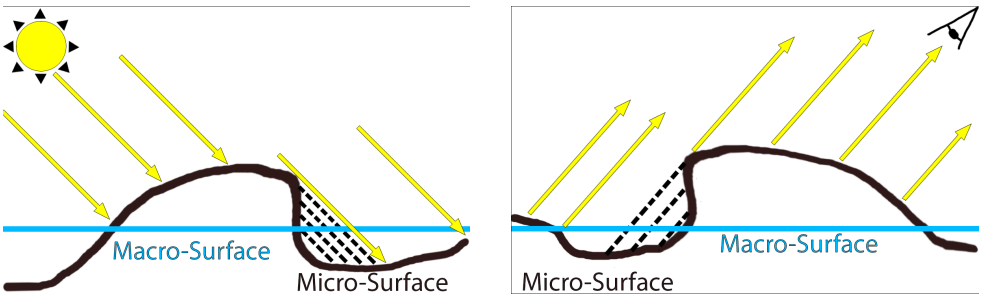


(a) Shadowing (b) Masking

Figure 2.8: (a) Due to the microgeometry, some microfacets are occluded and do not receive light (shadowing). (b) The light reflected from microfacets not visible from the viewing direction can not be seen (masking)

by accurate and adjustable formulae, however the most common mathematical model has the form:

$$f_r(\mathbf{v_i}, \mathbf{v_r}) = \frac{D \cdot G \cdot \Im}{4 \cos\theta_i \theta_r} \tag{2.5.1}$$

which also takes into account the Fresnel term $\Im$. Effect like masking and self-shadowing (see Figure 2.8) [62] depend on the projected area of the microfacets and hence on the distribution $D$, generally described by the geometrical attenuation term $G$; for a review of common masking functions and a derivation of the exact form of the masking function from the microsurface profile, see the work by Heitz [63]. This class of models can represent unique properties of the material and may include subsurface structure, generally resulting in complex calculations due to the interaction of the light with the surface structure.

### 2.5.1 Physically based models for Isotropic materials

The Cook-Torrance model [60] takes into account both specular and diffuse reflections, the latter modeled as Lambertian reflections. As for the specular component, the model assumes that only the fraction of the facets oriented in the direction of $h$ contributes to the final reflection, moreover it accounts for how many facets are visible from different view angles and how they reflect light [64]. These factors are modeled respectively through the functions $D$, $G$ and $\Im$:

$$f_{r,s}(\mathbf{v_i}, \mathbf{v_r}) = \frac{\Im(\theta_r) D(\mathbf{h}) G(\mathbf{v_i}, \mathbf{v_r})}{\pi \cos(\theta_r) \cos(\theta_i)}. \tag{2.5.2}$$

The expression of the distribution $D(\mathbf{h})$ is generally a Gaussian: $D(\mathbf{h}) = \cos(\theta_r)\exp^{-\left(\frac{\alpha}{m}\right)^2}$, where $\alpha$ is the angle between $\mathbf{v_i}$ and the reflected $\mathbf{v_r}$ and $m$ is a roughness parameter. The attenuation term $G$ includes both the shadowing and masking effects:

$$G(\mathbf{v_i}, \mathbf{v_r}) = \min\left(1, \frac{2(\mathbf{n}\cdot\mathbf{h})(\mathbf{n}\cdot\mathbf{v_r})}{\mathbf{v_r}\cdot\mathbf{h}}, \frac{2(\mathbf{n}\cdot\mathbf{h})(\mathbf{n}\cdot\mathbf{v_i})}{\mathbf{v_r}\cdot\mathbf{h}}\right). \tag{2.5.3}$$

One of the important contributions of this work is the formulation of the Fresnel term $\Im$, which represents the reflection of polished microfacets, approximated with the following expression:

$$\Im(\theta) = \frac{(g-c)^2}{2(g+c)^2}\left(1 + \frac{(c(g+c)-1)^2}{(c(g-c)+1)^2}\right) \tag{2.5.4}$$

where $c = \mathbf{v_r}\cdot\mathbf{h}$ and $g = \eta^2 + c^2 - 1$, being $\eta$ the index of refraction. The Cook-Torrance model can properly model metals, plastic with varying roughness and view-dependent changes in colour, although it does not follow the energy conservation principle in the entire hemisphere; additional drawbacks are the not intuitive parameters.

The He *et al.* model [65] accounts for polarisation and masking/shadowing effects, includes specular reflection when the surface roughness is low and takes into account the nature of light as an electromagnetic wave to model diffraction and interference, thus resulting into a very complex BRDF model. The model is able to represent metal, non-metal and plastic with smooth and rough surfaces and all parameters are physically based. The contribution to the reflection is given by three components, namely the specular term for mirror-like reflections, the directional diffuse and the uniform diffuse. The specular term describes mirror-like reflections from the mean plane of the surface:

$$f_{r,s}(\mathbf{v_i}, \mathbf{v_r}) = \frac{|F(\theta_r)|^2 \exp(-g(\sigma,\lambda))S(\mathbf{v_i}, \mathbf{v_r})}{\cos(\theta_i)d\omega_i}\Delta \tag{2.5.5}$$

where $F$ is the Fresnel reflectivity, $S$ is a shadowing function, $\sigma$ refers to the surface roughness, $\lambda$ is the wavelength of the incident light, $\Delta$ is a delta dirac function equal to 1 in the specular cone of reflection, and $d\omega_i$ is the incident solid

angle. The function of the surface roughness $g(\sigma,\lambda)$ is given by the expression $g(\sigma,\lambda) = (((2\pi\lambda)/\sigma)(\cos(\theta_i)+\cos(\theta_r)))^2$. For a smooth surface $S \to 1$ and $g \to 0$, hence the expression of the specular term becomes the specular reflectivity of a specular surface. As for the diffuse directional term, it describes diffraction and interference effects, which spread out the reflected field over the hemisphere, with a possible directional and nonuniform shape of the light intensity distribution:

$$f_{r,dd}(\mathbf{v_i},\mathbf{v_r}) = \frac{\Im(\mathbf{b},\mathbf{p})S\tau^2}{\cos(\theta_r)\cos(\theta_i)16\pi}\sum_{m=1}^{+\infty}\frac{g^m e^{-g(\sigma,\lambda)}}{m!m}e^{\left(-\frac{\mathbf{w_v}^2\tau}{4m}\right)}. \tag{2.5.6}$$

The directional diffuse reflection depends on surface roughness $\sigma$ and on the autocorrelation length $\tau$. The other parameters are the bisecting unit vector $\mathbf{b}$, the incident polarisation state vector $\mathbf{p}$ and the wave vector change $\mathbf{w_v}$. For very smooth surfaces $f_{r,dd}$ decreases to zero and for slightly rough surfaces the maximal values are aligned with the specular direction. As the roughness is increased the maximal values progressively move from off-specular angles to grazing angles for very rough surfaces. The uniform diffuse term is approximated with a Lambertian model and denoted by $f_{r,ud}(\mathbf{v_i},\mathbf{v_r}) = a(\lambda)$. An experimental analysis reported in [66] indicates that when polarisation and spectral dependencies are omitted, the He *et al.* model does not produce noticeably better visual results to the Cook-Torrance model [60]. The model does not suggest a sampling method and does not describe anisotropic materials.

Oren-Nayar [67] enhanced the Lambertian model for rough diffuse surfaces, to describe in a more realistic way the behaviour of real-world materials like concrete, sand and cloth, which show increasing brightness as the viewing direction approaches the light source direction, rather than being independent of the viewing direction. A rough diffuse surface is modelled as a collection of long symmetric V-cavities, each of which consists of two microfacets with a Lambertian reflectance; microfacets orientated toward the light source diffusely reflect some light back to the light source (backscatter). The model takes into account masking, shadowing and inter-reflections. The expression is given by:

$$f_r(\mathbf{v_i},\mathbf{v_r}) = \frac{\rho_d}{\pi}(A + B\max(0,\cos(\phi_i-\phi_r))\sin(\alpha)\tan(\beta)) \tag{2.5.7}$$

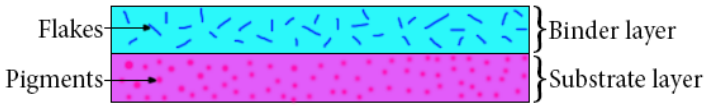


Figure 2.9: The bi-layered model by Ershov *et al.* [3]. The substrate layer is a solid paint film where the reflectance is Lambertian and the transparent binder layer contains embedded flakes.

where $\alpha = \max(\theta_r, \theta_i)$; $\beta = \min(\theta_r, \theta_i)$; given the surface roughness $\sigma$, the expressions for $A$ and $B$ are:
$A = 1 - [(0.5 \cdot \sigma^2)/(\sigma^2 + 0.33)]$; $B = (0.45 \cdot \sigma^2)/(\sigma^2 + 0.09)$. This model, widely used in Computer Graphics, obeys the reciprocity principle and reduces to the Lambertian model when $\sigma = 0$.

The multilayered model by Ershov *et al.* [68] represents car paint and consists of binder pigment particles, flakes and flake coatings. The model approximates the BRDF of each sub-layer and then merges sub-layers together and it is able to produce realistic appearance for car paints and models their components (binder, pigment particles, flakes). However, due to the complexity of the layered model, the computational time is significantly high. An updated version of the model is simplified to a bi-layered model and presents a substrate layer as a solid paint film where the reflectance is Lambertian and a transparent binder layer with embedded flakes (see Figure 2.9). Flakes are considered as partially transparent coloured mirrors, with the assumption that the reflectance of flakes does not depend on the incident direction and inter-reflections between flakes, so their interaction with light is modelled using constant reflectance values [3]. A large number of parameters are required and not all of them can be directly measured. This model is suitable for interactive design of automotive paints, by solving through optimisation the problem of finding pigment composition of a paint from its bidirectional reflectance distribution function.

Walter *et al.* [69] extend the microfacets theory introduced by [60] to simulate transmission through etched glass and other rough surfaces, thus taking into account the BSDF. The work by Smith [70], which investigated the geometrical self-shadowing of a surface described by Gaussian statistics, is also extended by deriving a shadowing function from any microfacet distribution $D$; the BRDF component follows 2.5.1. The distribution $D$ is different from previous models

and has been developed to better fit measured data; it is named GGX and has the following expression:

$$D(\mathbf{h}) = \frac{{\alpha_g}^2 \chi^+(\mathbf{h}\cdot\mathbf{n})}{\pi \cos^4\theta_h \left({\alpha_g}^2 + \tan^2\theta_h\right)^2} \tag{2.5.8}$$

where ${\alpha_g}^2$ is a width parameter and $\chi^+(x)$ is equal to one if $x > 0$ and zero if $x \leq 0$. The GGX distribution has a stronger tail than commonly used distributions, such as Beckmann and Phong, and thus tends to have more shadowing; in [71] it has been observed that the GGX distribution is identical to the Trowbridge-Reitz distribution [72]. From $D$ it is possible to derive a simple sampling equation and the expression of $G$, which is given by:

$$G(\mathbf{v_i}, \mathbf{v_r}, \mathbf{h}) \approx G_1(\mathbf{v_i}, \mathbf{h}) G_1(\mathbf{v_r}, \mathbf{h}) \tag{2.5.9}$$

$$G_1(\mathbf{v_x}, \mathbf{h}) = \chi^+\left(\frac{\mathbf{v_x}, \mathbf{h}}{\mathbf{v_x}, \mathbf{n}}\right) \frac{2}{1 + \sqrt{1 + {\alpha_g}^2 \tan^2\theta_x}}. \tag{2.5.10}$$

As previously observed, the GGX distribution has a longer tail than other distributions, however the GGX distribution fails to properly capture the glowy highlights of highly polished surfaces like the chrome sample in the MERL database [4], with a narrow specular peak and a much wider specular tail [73]. An anisotropic extension of the distribution, named Generalised-Trowbridge-Reitz, has been proposed by Burley [73]; a symmetric extension of the GGX to the entire ellipsoid domain, suitable for volumetric anisotropic materials, is described by Heitz *et al.* [74].

Another method aimed to describe the complex reflectance behaviour of a car paint is described in Rump *et al.* [75], which represents the reflectance with the first hybrid analytical BRDF and image-based BTF representation; the acquisition setup is described in Section 3.5. The appearance of metallic car paint is separated into the homogeneous BRDF part, which describes the reflection behaviour of the base and the top layer of the paint, and the spatially varying BTF part, which is caused by the aluminium flakes. The homogeneous part is represented by a multi lobe version of the Cook Torrance model [60]. In order to account for the characteristics of pearlescent paint, which show view-dependent off-specular colour

changes, the model includes a spectral view and light dependent part. The BRDF parameters are derived from the BTF measurements by means of a fitting procedure; the BRDF is calculated for every pixel and subtracted in the RGB space from the captured images. The resulting images contain only flakes data and they are used for a copy and paste synthesis approach.

Kurt *et al.* [76] proposed a BRDF model based on the halfway vector representation and Beckmann distribution. The model is physically plausible, can represent anisotropic materials, can accurately fit data and suggests an efficient importance sampling method, based on the strategy proposed by Ward *et al.* but with a different weighting function, which makes it particularly suitable for Monte Carlo Rendering algorithms. The basic BRDF model they propose is the sum of a pure Lambertian term and a single specular lobe, which can be readily extended to multiple specular lobes representation, to model mixture materials like a car paint:

$$f_r(\mathbf{v_i}, \mathbf{v_r}) = \frac{k_d}{\pi} + \sum_{l=1}^{N} \frac{k_{sl} \Im_l(\mathbf{v_r}, \mathbf{h}) D_l(\mathbf{h})}{4(\mathbf{v_r} \cdot \mathbf{h}) \left((\mathbf{v_i} \cdot \mathbf{n})(\mathbf{v_r} \cdot \mathbf{n})\right)^{\alpha_l}} \tag{2.5.11}$$

where $N$ is the number of lobes, $k_d$ is the diffuse albedo, $k_{sl}$ is the specular reflectivity per-lobe, $\Im_l$ is a per-lobe Fresnel term, $D_l$ a per-lobe normalised microfacet distribution, $\alpha_l$ is a set of parameters which needs to be chosen carefully to enforce energy conservation.

Bagher *et al.* suggested a function of $\tan^2 \theta_h{}^{-p}$ for the distribution $D$, where $p$ depends on the model [71], in order to enhance data fitting for single-layered materials like metals, metallic paints and shiny plastics, otherwise very difficult to fit with commonly used distributions and generally requiring several lobes, due to the shape of the decrease in their BRDFs, close to exponential at large angles but sharper at small angles. The model presented is the Cook−Torrance [60], in which the microfacets distribution is designed to efficiently and accurately approximate measured data. The distribution resulting from the suggested slope is called SGD (Shifted Gamma Distribution):

$$D(\theta_h) = \frac{\chi_{[0,\pi/2]}(\theta_h) \alpha^{p-1} e^{-\frac{\alpha^2 + \tan^2 \theta_h}{\alpha}}}{\pi \cos^4 \theta_h \Gamma(1-p, \alpha) \left(\alpha^2 + \tan^2 \theta_h\right)^p} \tag{2.5.12}$$

where $\alpha$ is a fitting parameter, $\chi_{[0,\pi/2]}(\theta_h)$ is equal to 1 if $\theta_h < \pi/2$ and 0 otherwise, $\Gamma$ is the incomplete Gamma function:

$$\Gamma(1-p,\alpha) = \int_{\alpha}^{\infty} t^{-p} e^{-t} dt. \tag{2.5.13}$$

From the SGD it is possible to derive the shadowing function $G$ and a sampling method.

Low *et al.* [77] proposed two isotropic models for glossy surfaces, based either on the Rayleigh-Rice light scattering theory (smooth surface BRDF) or on the microfacet theory (microfacet BRDF). Both models make use of a modified version of the ABC model [78], [79], which was originally formulated to fit the Power Spectral Density of some measured smooth surfaces. The PSD describes the surface statistics in terms of the spacial frequencies $f_x$ and $f_y$, which depend on the wavelength $\lambda$ of the incident light:

$$f_x(\mathbf{v_i}, \mathbf{v_r}) = (\sin\theta_r \cos\phi_r - \sin\theta_i)/\lambda; \tag{2.5.14}$$

$$f_y(\mathbf{v_i}, \mathbf{v_r}) = (\sin\theta_r \sin\phi_r)/\lambda. \tag{2.5.15}$$

The ABC model [78], [79] is able to model the inverse power law shape PSD of polished data, and it is given by:

$$PSD(f) = A^{'} / \left(1 + B^2 f^2\right)^{\frac{C+1}{2}} \tag{2.5.16}$$

where A is determined by low-frequency spectral density, $B = 2\pi l_0$, $l_0$ is the autocorrelation length, $C > 0$, $f = \sqrt{{f_x}^2 + {f_y}^2}$, $A^{'} = \Gamma((c+1)/2)AB/[2\Gamma(c/2)\sqrt{\pi}]$ and $\Gamma$ is the gamma function. In [77] the ABC model is simplified to $S(f) = a/\left(1+bf^2\right)^c$, where the mapping of the new parameters to the original ABC is: $a = A^{'}$, $b = B^2$ and $c = (C+1)/2$; in practice narrower specular peaks are obtained by increasing $b$, whereas $c$ controls the fall-off rate of wide-angle scattering. The smooth surface BRDF has the following expression:

$$f_r(\mathbf{v_i}, \mathbf{v_r}) = (k_d/\pi) + O\,\Im(\theta_d)\, S(\|\mathbf{D_p}\|) \tag{2.5.17}$$

where $k_d$ is a scaling factor for the Lambertian term, $O$ is a modified obliquity fac-

tor, $\Im(\theta_d)$ is the Fresnel term in Equation 2.5.2 with extinction coefficient set to zero and aimed to approximate the reflectivity polarisation factor, which depends on the surface material properties. $\mathbf{D_p}$ is the projected deviation vector, defined as $\mathbf{D_p} = \mathbf{v_{r,p}} - \mathbf{r_{v_{i,p}}}$, where $\mathbf{v_{r,p}}$ is the projection of $\mathbf{v_r}$ on the surface tangent plane and $\mathbf{r_{v_{ip}}}$ is the projection of the mirror direction of $\mathbf{v_i}$ on the surface tangent plane. To deal with unreliable data near grazing angles, the value suggested for the obliquity factor $O$ is 1 instead of the typical definition of $O = \cos\theta_i \cos\theta_r$. The microfacet model is based on Cook-Torrance [60] and makes use of the modified ABC distribution:

$$f_r(\mathbf{v_i}, \mathbf{v_r}) = \frac{k_d}{\pi} + \frac{\Im(\theta_h) S\left(\sqrt{1 - \mathbf{h}\cdot\mathbf{n}}\right) G(\mathbf{v_i}, \mathbf{v_r})}{\mathbf{v_i}\cdot\mathbf{n}\ \mathbf{v_r}\cdot\mathbf{n}} \tag{2.5.18}$$

where $\Im$ and $G$ are the same as in Equation 2.5.2, $S$ is the modified ABC distribution and $k_d$ is again a scaling factor for the diffuse component; the parameter $a$ of $S$ is used as a scaling factor for the specular term, hence the distribution is not normalised. The model is reciprocal but does not obey energy conservation. Both models provide accurate fits to measured data, with the microfacet model showing lower errors, and accurately represent scattering from glossy surfaces with sharp specular peaks and non-Lambertian wide angle scattering. For both models an efficient importance sampling strategy is suggested.

The discrete stochastic model by Jakob *et al.* [80] extends the microfacet theory by replacing the continuous distribution of microfacets in the Cook-Torrance model [60] with a discrete one, thus assuming that a surface consists of a high but finite number of scattering particles. This assumption facilitates modelling a controllable, non-smooth spatially varying BRDF appearance of a glittery surface, like mica flakes, ice crystals, metallic car paint and craft glitter for decorations. The notion of multiscale BRDF is introduced, which takes into account finite areas and solid angles rather than single points and directions:

$$f_r(A, \mathbf{v_i}, \omega_r) = \frac{(\mathbf{v_i}\cdot\mathbf{h})\Im(\mathbf{v_i}\cdot\mathbf{h}) D(A, \omega_h), G(\mathbf{v_i}, \mathbf{v_r}, \mathbf{n})}{a(A)\sigma(\omega_r)(\mathbf{v_i}\cdot\mathbf{n})(\mathbf{v_r}\cdot\mathbf{n})} \tag{2.5.19}$$

where $A$ is the area around the point $p$ into account, $a(A)$ its surface area, $\omega_h := \{(\mathbf{v_i} + \mathbf{v_r}/\|\mathbf{v_i} + \mathbf{v_r}\|), \mathbf{v_r} \in \omega_r\}$ is the set of microfacet normals that reflect from $\mathbf{v_i}$ into the finite solid angle $\omega_r$ around $\mathbf{v_r}$, $\sigma(\omega_r)$ is the area of $\omega_r$ on the unit sphere,

$\Im$ is the fresnel term, $G$ models shadowing and masking. The discrete multiscale microfacets distribution $D$ is defined as:

$$D(A,\omega_h) = \frac{1}{N}\sum_{k=1}^{N} 1_{\omega_h}(\mathbf{v_h}^k)1_A(p^k) \tag{2.5.20}$$

where $p^k$ and $\mathbf{v_h}^k$ are the position and normal of the $k^{th}$ microfacet of a list of $N$ microfacets, $1_A$ and $1_{\omega_h}$ are the indicator functions of the sets $A$ and $\omega_h$ respectively. The indicator functions control the appearance of the surface, since they determine which microfacets in $A$ reflect light into the solid angle $\omega_r$ around $\mathbf{v_r}$: a high number of participating facets gives a smoother appearance than a low number, which gives instead a strongly glittery appearance. An efficient implementation of the model is discussed, together with an importance sampling strategy for Monte Carlo renderers.

### 2.5.2 Physically based models for Anisotropic materials

To simulate both smooth and rough multi-layered materials, Weidlich and Wilkie [81] proposed to combine several microfacet based layers into a single physically plausible BRDF model. Their model assumes that any microfacet is large in relation to the layer thickness, models the absorbtion of part of light when it travels inside a transparent material and include a total reflection term, when light propagates at an angle of incidence greater than the critical angle; the simplicity of the model does not allow reproducing effects like iridescence.

Dupuy *et al.* [82] proposed an approach to automatically convert an arbitrary material to a microfacet BRDF. The facet distribution is obtained by solving an eigenvector problem, based solely on backscattering samples and simplifying the Fresnel term to a constant; once an eigenvector with all positive components is found with the power iteration method, its values are linearly interpolated to build a continuous distribution. The Fresnel term is then recovered by calculating for each colour channel the average ratio between the input and an ideal mirror microfacet BRDF with a constant Fresnel term equal to 1, in the form $f_r(\mathbf{v_i},\mathbf{v_r}) = D(\mathbf{h})G(\mathbf{v_i},\mathbf{v_r})/(4\cos\theta_i\cos\theta_r)$. This method can fit an anisotropic BRDF in a few seconds and allows to edit the properties of the rough-

ness distribution, however its accuracy depends on the density of the measurements in the backscattering direction, thus limiting the applicability to more complex anisotropic BRDFs [83].

## 2.6 Data-Driven models

Data-Driven models approximates measured BRDFs with a suitable function space, e.g. spherical harmonics or wavelets, weighted sum of separable functions or product of functions. Measured BRDF data, produced by most of the setups described in Chapter 3, can be stored in a table or a grid and then interpolated, to produce a large look-up table when data is needed. This method is simple but inefficient in terms of storage. Moreover, the measured raw data is often noisy, hence the noise is likely to appear in the rendered material. A measured BRDF can be fitted to analytic models and employed to reconstruct the BRDF, thus significantly reducing storage size. The down side of this strategy is related to the inflexibility of many models, hard to edit and able to represent only limited classes of materials. A different solution is to approximate measured BRDFs with a suitable function space, *e.g.* spherical harmonics or wavelets, weighted sum of separable functions or product of functions. We refer to this class of models as Data-Driven models.

The general idea behind this class of models use the measured data to derive some ad-hoc model, generally based on the principle that a continuous function can be represented by a linear combination of basis functions and a mixture of basis functions can be used for interpolation.

A continuous function can be represented by a linear combination of basis functions and a mixture of basis functions can be used for interpolation. In the Fourier basis, the functions are expressed as a sum of sinusoidal and cosinusoidal terms. In the polynomial basis a collection of quadratic polynomials are used with real coefficients.

A possible way to represent BRDFs is to project them onto an orthonormal basis [62], mapped onto a unit disc and projected on to a hemisphere. Spherical wavelets, spherical harmonics and Zernike polynomials are mathematically and computationally efficient, since they represent the shape of the BRDF as

the sum of low and high frequency functions to capture the shape of the BRDF. Wavelets [84–86] can represent large specular peaks more efficiently than spherical harmonics [87] and Zernike polynomials [27]. The limitation of these methods is the significant memory requirement even to obtain simple BRDFs, since a large number of basis functions is generally required.

Separable decompositions of a high-dimensional function $f$ can be used to approximate it to arbitrary accuracy, using a sum of products of lower-dimensional functions. Four dimensional BRDFs can be written as a sum of terms each of which is the product of two-dimensional functions $g(\cdot,\cdot)$ and $h(\cdot,\cdot)$:

$$f_r(\mathbf{v_i},\mathbf{v_r}) = f_r(\theta_i,\phi_i,\theta_r,\phi_r) \approx \sum_{k=1}^{N} g_h(\theta_i,\phi_i)h_k(\theta_r,\phi_r). \tag{2.6.1}$$

This representation directly approximates the fully tabulated BRDF over all directions, implementing manageability of data for use in rendering systems and it is also useful for the purpose of importance sampling. If a good approximation can be found for a small $N$, a separable decomposition is capable of high compression rates, thus resulting in a compact way to store large measured datasets, while maintaining accurate representation. The parameterisation of the lower-dimensional functions can improve the performance of the decomposition and needs to be wisely chosen in order to minimise the number of functions needed for BRDF representations. A common reparameterisation makes use of the angle halfway the incident and exitant directions and the difference angle [27] (see Figure 2.1(b)):

$$f_r(\mathbf{v_i},\mathbf{v_r}) \approx \sum_{k=1}^{n} g_k(\mathbf{v_h})h_k(\mathbf{v_d}). \tag{2.6.2}$$

where $\mathbf{v_h}$ and $\mathbf{v_d}$ arise from the re-parameterisation.

A common technique to obtain a separable representation is the Singular Value Decomposition (SVD) [88]. Given a matrix $M$ its SVD is the factorisation in the form $M = USW^T$, where $S = diag(\sigma_k)$ is a diagonal matrix of singular values; the columns of $U = [\mathbf{u}_k]$ and $W = [\mathbf{w}_k]$ are orthonormal. The matrix $M$ can be written as:

$$M = USW^T = \sum_{k=1}^{K} \sigma_k \mathbf{u}_k {\mathbf{w}^T}_k \tag{2.6.3}$$

where $\mathbf{u}_k \mathbf{w}^T{}_k$ is an outer product. In Fournier [89] the SVD decomposition is used to approximate the Blinn-Phong BRDF using Ward's measured data; the BRDF is approximated with a sum of terms each of which is the product of two functions, one of the incident and one of the outgoing direction.

### 2.6.1 Data-driven models for Isotropic materials

Matusik *et al.* [4] presented a set of data-driven reflectance models, based on either linear and non-linear dimensionality reduction (Slides 48-49). A set of 104 isotropic BRDFs, parameterised using the half-angle [27], are discretised into $90 \times 90 \times 180$ bins which are smoothed by removing outliers. The linear dimensionality reduction used is the Principal Component Analysis (PCA), which allows determination of a set of basis vectors that span the linear subspace on which the BRDFs lie. The RGB colour channels are assembled together and analysed in the log space, in order to reduce the difference between specular and non-specular values. A linear combination of a subset of the principal components is used for the reconstruction and in most cases 30-40 components give good results. As for the non-linear dimensionality reduction, the charting algorithm [90] has been used, since it gives good results even with a small number of samples and at the same time it reduces the noise in the data. The idea behind charting is that data lies on a low dimensional manifold embedded in the sample space and tries to find a kernel-based mixture of linear projections to smoothly map the samples on the coordinate system, while preserving local relationships between the sample points. Since each dimension performs a noise suppression in a different direction the error does not decrease monotonically. However they found that the BRDF data lies in a 10D manifold and a 15D manifold would suffice to synthesise new BRDFs even over long distances. The advantages of such data-driven BRDF models is the realistic appearance and meaningful parameterisation.

Matusik *et al.* performed a wavelet analysis for all of the measured isotropic BRDFs [91] (Slide 49), in order to find the maximum required frequency to sample any arbitrary BRDF correctly. For each BRDF a non-uniform wavelet transform is applied to determine the highest coefficients able to reconstruct the BRDF itself with high precision, while setting to zero the rest of the coefficients. The

union of the sets of non-zero wavelet coefficients (about 69,000), which show some degree of coherence among different BRDFs, corresponds to a set of wavelet functions called Common Wavelet Basis (CWB). The CWB allows reconstruction of a BRDF by solving a system of linear equations. In the same work a simple approach to represent a new measured BRDF is presented, using a linear combination of the BRDFs in the dataset. Using this data they construct a over-constrained system of equation in the form $P \times C \approx B$, where $P$ is the matrix of the BRDFs in the dataset, $C$ a vector of coefficients and $B$ the new measured BRDF. A subset $X$ of the rows of $P$ is constructed in such a way the ratio between the highest and lowest eigenvalue of the matrix $X^T X$ is small. Experimentally they have found that 800 samples are enough to represent a new BRDF.

Romeiro *et al.* [92] describe a method for inferring the reflectance of isotropic materials from images, assuming known curved surface with known natural illumination. To reduce the dimension of the BRDF domain, parameterised using the halfway vector and difference angle [27], thanks to the reciprocity assumption it is possible to apply the projection $\phi_d \leftarrow \phi_d + \pi$; for isotropic materials it is possible to apply the projection onto the domain $(\theta_h, \theta_d, \phi_d)$ and in case of bilateral symmetry (*i.e.* if the reflectance of the material shows little changes when $\mathbf{v_r}$ is reflected about the incident plane) it is possible to apply the projection $\phi_d \leftarrow \phi_d + \pi/2$. If a material is bivariate, *.i.e.* it satisfies a further generalisation of isotropy, bilateral symmetry and reciprocity, the projection onto the domain $(\theta_h, \theta_d) \in [0, \pi/2]$ is allowed. A bivariate representation is often sufficient to capture off-specular reflections, retro-reflection and Fresnel effect. Under these assumptions, the resulting 2D domain is sampled using the functions $s(\theta_h, \theta_d) = 2\theta_d/\pi$ and $t(\theta_h, \theta_d) = \sqrt{2\theta_h/pi}$, which allow to increase the sampling density near specular reflections; since bivariate BRDFs vary slowly over a significant region of their domain, an optimisation framework with a smoothness constraint is employed to recover the BRDF:

$$\underset{f_r \geq 0}{\operatorname{argmin}} \parallel I - L f_r \parallel_2^2 + \alpha \left( \parallel \Lambda_s^{-1} D_s f_r \parallel_2^2 + \parallel \Lambda_t^{-1} D_t f_r \parallel_2^2 \right) \tag{2.6.4}$$

where $\mathscr{S} = (s_i, t_i)$ is a uniform grid in the BRDF domain, $D_s$ and $D_t$ are $|\mathscr{S}| \times |\mathscr{S}|$ derivative matrices, $\alpha$ is a regularisation parameter, $\Lambda_s$ and $\Lambda_t$ are $|\mathscr{S}| \times |\mathscr{S}|$

matrices that control non-uniform regularisation in the $(\theta_h, \theta_d)$ domain and $L$ is a lighting matrix. The term $I$ is related to the rendering equation:

$$I(\mathbf{v_r}, \mathbf{n}) = \int_\Omega L(R_n^{-1}\mathbf{v_i}) f_r(s(\mathbf{v_i}, R_n\mathbf{v_r}), t(\mathbf{v_i}, R_n\mathbf{v_r})) \cos\theta_i dv_i \tag{2.6.5}$$

where $R_n$ rotates $\mathbf{n}$ towards the $z$-axis and $\mathbf{v_r}$ towards the $xz$-plane. In order to obtain good results, the environment illumination used to capture the 2D picture must allow sufficient observations of the BRDF $(\theta_h, \theta_d)$ domain, in particular in regions corresponding to specular reflections, retro-reflections and grazing angles, which occur respectively at $(\theta_h \approx 0)$, $(\theta_d \approx 0)$ and $(\theta_d \approx \pi/2)$.

### 2.6.2 Data-driven models for Anisotropic materials

A technique for separable decomposition of BRDFs based on either Singular Value Decomposition (SVD) (eq. 2.6.8) or Normalised Decomposition (ND) (eq. 2.6.9) is described by Kautz and McCool [93]. In both cases the separable decomposition $f_{sd,r}$ to approximate the BRDF $f_r$ has the form:

$$f_r(\mathbf{v_i}, \mathbf{v_r}) = f_{sd,r}(\mathbf{P_x}(\mathbf{v_i}, \mathbf{v_r}), \mathbf{P_y}(\mathbf{v_i}, \mathbf{v_r})) \tag{2.6.6}$$

where $P_x$ and $P_y$ are vector functions. In the following, the parameters of $f_{sd,r}$ are $\mathbf{x} = \mathbf{P_x}(\mathbf{v_i}, \mathbf{v_r})$ and $\mathbf{y} = \mathbf{P_y}(\mathbf{v_i}, \mathbf{v_r})$. The matrix $M$ of Equation 2.6.3 consists of the tabulated and reparameterised BRDF values of $f_{sd,r}(\mathbf{x}, \mathbf{y})$:

$$M = \begin{pmatrix} f_{sd,r}(\mathbf{x}_1, \mathbf{y}_1) & \cdots & f_{sd,r}(\mathbf{x}_1, \mathbf{y}_K) \\ \vdots & \ddots & \vdots \\ f_{sd,r}(\mathbf{x}_K, \mathbf{y}_1) & \cdots & f_{sd,r}(\mathbf{x}_K, \mathbf{y}_K) \end{pmatrix} \tag{2.6.7}$$

The resulting $\mathbf{u}_k$ and $\mathbf{w}_k$ from the SVD of $M$ can be interpolated in order to obtain the 2D functions $u_k(\mathbf{x})$ and $w_k(\mathbf{y})$:

$$f_{SVD,r}(\mathbf{x}, \mathbf{y}) \approx \sum_{k=1}^{N} \sigma_k u_k(\mathbf{x}) w_k(\mathbf{y}) \tag{2.6.8}$$

As for the ND factorisation:

$$f_{ND,r}(\mathbf{x},\mathbf{y}) \approx g_1(\mathbf{x},\mathbf{y})\, h_1(\mathbf{x},\mathbf{y}) \tag{2.6.9}$$

where $g_1$ is a constant if $\mathbf{P}_\mathbf{x}(\mathbf{v_i},\mathbf{v_r})$ is fixed and scales the profile $h_1$. While SVD can produce optimal approximations and minimises the RMS error, it is expensive in terms of time and space resources and can produce negative factors in the expansion; the ND algorithm does not guarantee optimality but requires less memory than the SVD decomposition and it is faster. The lower dimensional functions are stored into texture maps, to allow multiplications being performed by compositing or multitexturing.

McCool and Ahmad presented [94] a decomposition algorithm for both isotropic and anisotropic BRDFs. The algorithm is based on logarithmic homomorphism (eq. 2.6.10) and it is general enough to approximate BRDFs with an arbitrary number of positive factors and degree of precision, while satisfying the Helmholtz reciprocity, but limited to point and directional light sources. The authors describe a simple parameterisation (eq. 2.6.11) and demonstrate that it is possible to limit the storage cost to just two texture maps, obtaining good compression ratios:

$$\log\left(f_r(\mathbf{v_i},\mathbf{v_r})\right) \approx \sum_{j=1}^{N} \log\left(p_j\left(\pi_j(\mathbf{v_i},\mathbf{v_r})\right)\right), \tag{2.6.10}$$

$$f_r(\mathbf{v_i},\mathbf{v_r}) \approx p(\mathbf{v_i})q(\mathbf{h})p(\mathbf{v_r}) \tag{2.6.11}$$

where $p(\cdot)$ are two dimensional functions and $\pi_j$ are projection functions $\mathbb{R}^4 \rightarrow \mathbb{R}^2$. The logarithmic transformation tends to disregard large peaks in the data and smoothes specular highlights, which may lead to high approximation errors.

Lawrence *et al.* [41] presented an importance sampling algorithm for arbitrary BRDFs, based on reparameterising the BRDFs using the half-angle or the incident angle, followed by a non-negative matrix factorisation, essential for sampling purposes:

$$f_r(\mathbf{v_i},\mathbf{v_r})(\mathbf{v_i}\cdot\mathbf{n}) \approx \sum_{j=1}^{J} F_j(v_r) \sum_{k=1}^{K} u_{jk}(\theta_p) v_{jk}(\phi_p) \tag{2.6.12}$$

The factored form (Equation 2.6.12) allows expressing the BRDF, multiplied by

the cosine of the incident angle, as a sum of a small number of terms, each of which is a product of a 2D function $F_j$ only dependent on the outgoing direction and two 1D functions $u_{jk}, v_{jk}$ dependent on the angle chosen for the parameterisation $\mathbf{v_p} = (\theta_p, \phi_p)$. The 1D functions are used to interpret the factors as 1D probability distributions. The results are generally accurate and the technique can be used for sampling BTFs and light fields, but does not enforce reciprocity and the representation may present a discontinuity at the pole of the angle selected for the parameterisation.

In later work Lawrence *et al.* [95] presented an algorithm based on linear constraint least squares, capable of compact and accurate SVBRDF representation for rendering. Under the assumption that BRDFs are blended linearly over the surface, the matrix factorisation algorithm provides an editable decomposition and can represent directional and spatial reflectance behaviour of a material. The described Inverse Shade Tree (IST) representation takes as input a measured materials dataset and a user-supplied tree structure and fills in the leaves of the tree. IST proceeds top-down at each stage decomposing the current dataset according to the type of node encountered in the tree. The leaves provide editability since they correspond to pieces that are meaningful to the user. The Alternating Constrained Least Squares algorithm (ACLS) decomposes the SVBRDF into basis BRDFs as 4D functions in tabular form which are then decomposed into 2D functions and further into 1D curves. The 1D curves represent data simply and accurately for isotropic materials, for anisotropic materials the decomposition ends into 2D functions. The advantage of ACLS is the possibility to easily add linear constraints, thus allowing to enforce energy conservation, reciprocity and monotonicity, other than sparsity and non-negativity, but it requires building a regularly sampled data matrix for factorisation.

Tensor representation has been previously used for interactive modification of the material properties and relighting by Sun *et al.* [96]. Based on the observation that high-frequency specular lobes generally require a large number of basis terms for reconstruction, thus precluding interactive performance, the BRDFs are separated into a specular lobe $f_{r,s}$ and the remainder $f_{r,rm}$. The specular lobe $f_{r,s}$ is modeled as a sum of 4 Gaussians, with different neighborhood support ($0^\circ$ for perfect mirror reflection, $7^\circ$, $14^\circ$ and $21^\circ$ for broader Gaussians). By removing

$f_{r,s}$ the BRDF is left with mainly low frequency terms, that can be modeled with a small basis by tensor approximation.

When a non-linear function is used for BRDF data fitting there are several shortcomings, due the number of parameters which can be large depending on the model and the number of lobes, and to the non-linear estimation process which can be computationally expensive. Ozturk *et al.* [56] proposed a representation based on response surface models, defined as a polynomial function of order $p$ in $k$ variables, and expressing a BRDFs as functions of the incoming and outgoing direction and transforming the variables of some non-linear reflectance models (specifically Ward [49], Lafortune [52] and Ashikhmin- Shirley [55], described in Section 2.4) using Principal Component Analysis, thus obtaining a linear representation. This reciprocal but not energy preserving representation is general enough to model both isotropic and anisotropic materials, diffuse and glossy.

Bilgili *et al.* [97] proposed to represent four-dimensional measured BRDFs data as a function of tensor products, factorised using Tucker decomposition [98], a generalisation of higher order principal component analysis. Tensors are a generalisation of scalars and vectors to higher orders and their rank is defined by the number of directions, *e.g.* a scalar is a zero-order tensor and a vector a first-order tensor; the Tucker factorization decomposes a tensor into a set of matrices and one small core tensor. The logarithmic transformation of a $4D$ BRDF data matrix $B = b_{ijkl}$, based on the halfway vector representation, can be roughly approximated by setting all the Tucker parameters to 1:

$$\log(b_{ijkl}) \approx g f_1(\theta_{hi}) f_2(\phi_{hj}) f_3(\theta_{rk}) f_4(\phi_{rl}) \tag{2.6.13}$$

where $i = 1,\ldots,N_{\theta_h}$, $j = 1,\ldots,N_{\phi_h}$, $k = 1,\ldots,N_{\theta_r}$, $l = 1,\ldots,N_{\phi_r}$ and $N_{\theta_h}$, $N_{\phi_h}$,$N_{\theta_r}$,$N_{\phi_r}$ are the sampling resolution of the BRDF data, $g$ is the zero-order core tensor, $f_1(\theta_{hi})$, $f_2(\phi_{hj})$, $f_3(\theta_{rk})$, $f_4(\phi_{rl})$ are univariate tensor functions respectively evaluated at $\theta_{hi}$, $\phi_{hj}$, $\theta_{rk}$ and $\phi_{rl}$; the logarithmic transformation eliminates the problem of estimated negative BRDF values. The error matrix $\mathbf{e}_1$ of this approximation can be written as $\mathbf{B}_0 = \mathbf{B}_0^{'} + \mathbf{e}_1$, where $\mathbf{B}_0 = \log(b_{ijkl})$ and $\mathbf{B}_0^{'}$ is the approximation. The approximation is improved by applying recursively the decomposition on the error terms, until a satisfactory level of accuracy is obtained; assuming that $S$ is

the total number of iterations, the expression of $\mathbf{B}_0$ becomes:

$$\mathbf{B}_0 \approx \mathbf{B}_0^{'} + \mathbf{e}_1^{'} + \mathbf{e}_2^{'} + \ldots + \mathbf{e}_{S-1}^{'} \tag{2.6.14}$$

where $\mathbf{e}_1 = \mathbf{e}_1^{'} + \mathbf{e}_2$ and $\mathbf{e}_1^{'}$ is the Tucker approximation of $\mathbf{e}_1$, $\mathbf{e}_2$ is the error term of the second and so on. This non-negative representation allows good compression ratios while being able to represent Fresnel effects and off-specularities, but does not satisfy reciprocity and energy conservation. As for the importance sampling, to limit the sampled region for isotropic materials it has been shown that most of the total variation is due to two components which corresponds to univariate functions of $\theta_h$ and $\theta_r$; a similar property is observed for the anisotropic material, where the main components are univariate functions of $\theta_h$ and $\phi_h$.

Pacanowski *et al.* [99] employ a subset of the halfway parameterisation [27] to project measured BRDFs on the two-dimensional space $(\theta_h, \theta_d)$ and approximates the projection by using Rational Functions (RF), since they are able to properly approximate the typical steep changes of specular lobes. A RF $r$ of a finite dimensional vector $\mathbf{x}$ of real variables is defined as:

$$r_{n,m}(\mathbf{x}) = \frac{\sum_{j=0}^{n} p_j b_j(\mathbf{x})}{\sum_{k=0}^{m} q_k b_k(\mathbf{x})} \tag{2.6.15}$$

where $p_j$ and $q_k$ are real numbers and $b_j(\mathbf{x})$, $b_k(\mathbf{x})$ are multivariate basic functions, for example multinomials. Given $t+1$ measured values $b_i$ located at a vector $\mathbf{x}_i$ and contained in the intervals $[\underline{b_i}, \overline{b_i}]$, the data fitting problem can be stated as finding a RF $r_{n,m}(\mathbf{x})$ with the smallest possible $n+m$, to interpolate the $t+1$ intervals with the additional constraints of non-negativity, monotonicity and symmetry: $\forall i = 0, \ldots, t \quad \underline{b_i} \leq r_{n,m}(\mathbf{x}_i) \leq \overline{b_i}$. The widths of the interpolation intervals are chosen in such a way that the renderings are visually satisfactory, while keeping the number of coefficients reasonably low; the solution is found by solving a quadratic programming problem. Isotropic BRDF data are approximated with a single RF, called Rational BRDF:

$$f_{r,s}(\mathbf{v_i}, \mathbf{v_r}) \approx f_{\theta_h, \theta_d} \approx r_{n,m}(\theta_h, \theta_d). \tag{2.6.16}$$

The anisotropic model is based on the observation that for some anisotropic materials like brushed metals, the variation of the reflected intensity, when the surface is rotated around the normal **n**, consists of a scaling factor applied to an average isotropic lobe:

$$f_{r,s}(\mathbf{v_i}, \mathbf{v_r}) \approx r^a_{n',m'}(\phi_h) r^i_{n,m}(\theta_h, \theta_d) \tag{2.6.17}$$

where $r^i_{n,m}(\theta_h, \theta_d)$ is a isotropic Rational BRDF and $r^a_{n',m'}(\phi_h)$ is a scaling factor to model anisotropic variations. The same approximation process applied to BRDFs can be applied to the inverse Cumulative Distribution Function (CDF) and used for importance sampling; the use of RF for BRDFs and CDFs allows to obtain a very small memory footprint.

More recently tensor representation has been used in [100, 101] to represent anisotropic materials with no assumption on the reflectance and scattering behaviour, particularly useful in presence of unusual scattering properties. The measured data is fitted to a series of radial basis functions in order to derive a continuous representation from the sparse input 4-D measurements. The incident and reflected hemispheres are projected onto disks and mapped over the unit square; the four dimensions given by the two squares define a rank-4 tensor, subdivided into a tensor tree for fast Monte Carlo sample generation. The tensor tree representation adaptively subdivides sharp peaks of the BRDF in different regions of the distribution, with an additional averaging step between incident and reflected direction to account for Helmoltz reciprocity.

## 2.7 Conclusion

We described the state of the art in BRDF representation. Each model is limited to a particular set of parameters which result in ability of the model to represent a specific material group. Even generalised models can not cover a broad range of materials nor variations of a material within one group. As a consequence, existing model assets can rarely be reused and material modelling involves a great deal of manual effort. The broad range of material models and complexity of the parameters requires from an artist an understanding of the underlying representation and materials microstructure and macrostructure. A physically accurate, consis-

| | Model and Ref | Reciprocity | Energy Cons. | Fresnel | Sampling | Short Description |
|---|---|---|---|---|---|---|
| | **PHENOMENOLOGICAL** | | | | | |
| ISOTROPIC | Phong [43] | ✗ | ✗ | ✗ | ✓ | Basic surface representation model based on the Cosine law. |
| ISOTROPIC | Blinn - Phong [47] | ✗ | ✗ | ✗ | ✓ | Based on [43], uses the halfway reflection direction for faster computation. |
| ISOTROPIC | Nishino and Lombardi [102] | ✓ | ✓ | ✗ | ✗ | Models BRDFs as a mixture of hemispherical distribution functions; small footprint. |
| ISOTROPIC | Brady et al. [48] | ✓ | ✗ | ✓ | ✗ | Framework for automatic learning of analytical models. Some of the properties are not guaranteed by the grammar. |
| ANISOTROPIC | Ward [49] | ✓ | ✗ | ✗ | ✓ | Versatile and cheap to compute. |
| ANISOTROPIC | Lafortune [52] | ✓ | ✓ | ✗ | ✓ | Generalisation of the cosine lobe model with multiple steerable lobes. |
| ANISOTROPIC | Neumann et al. [50] | ✓ | ✓ | ✓ | ✓ | Physically plausible formulation of the Phong, Blinn-Phong and Ward models. |
| ANISOTROPIC | Ashikhmin-Shirley [55] | ✗ | ✗ | ✓ | ✓ | Based on [43], includes anisotropic reflections for two-layered materials. |
| ANISOTROPIC | Edwards et al. [57] | ✗ | ✓ | ✗ | ✓ | Framework for transforming the halfway vector into different domains. |
| ANISOTROPIC | Ashikhmin-Premoze [58] | ✓ | ✓ | ✓ | ✓ | Combines analytic model with a data-driven distribution; accounts for backscattering. |
| | **PHYSICALLY-BASED** | | | | | |
| ISOTROPIC | Cook-Torrance [60] | ✓ | ✗ | ✓ | ✗ | It can model metals and plastics, view dependent changes. |
| ISOTROPIC | He et al. [65] | ✓ | ✓ | ✓ | ✗ | Enhances [60], allowing more general material representation. |
| ISOTROPIC | Oren-Nayar [67] | ✓ | ✓ | ✗ | ✓ | Enhance the Lambertian model for rough diffuse surfaces. |
| ISOTROPIC | Ershov et al. [68] | ✗ | ✓ | ✓ | ✗ | Focuses on layered materials, like metallic paint. It models binder pigment particles, flakes and flake coating. |
| ISOTROPIC | Walter et al. [69] | ✓ | ✓ | ✓ | ✓ | Defines the GGX distribution; based on a BSDF representation. |
| ISOTROPIC | Rump et al. [75] | ✗ | ✗ | ✓ | ✓ | Suitable for metallic paints, combines [60] for the base layer with BTF for top paint layer, including particles. |
| ISOTROPIC | Kurt et al. [76] | ✓ | ✓ | ✓ | ✓ | The multiple specular lobe model can represent layered or mixed materials. |
| ISOTROPIC | Bagher et al. [71] | ✓ | ✗ | ✓ | ✓ | Provides accurate fitting for materials in the MERL database |
| ISOTROPIC | Löw et al. [77] | ✓ | ✗ | ✓ | ✓ | Guarantees accurate fitting to measured data for glossy surfaces; descibes 2 models based on the ABC distribution |
| ISOTROPIC | Jakob et al. [80] | ✗ | ✗ | ✓ | ✓ | Allows modelling spatially varying BRDF appearance of glittery surfaces |
| ANIS. | Weidlich and Wilkie [81] | ✓ | ✓ | ✓ | ✓ | Multi-layered model which includes absorbtion and internal reflection. |
| ANIS. | Dupuy et al. [82] | ✗ | ✗ | ✓ | ✓ | Method to automatically convert a material to a microfacet BRDF. |
| | **DATA-DRIVEN** | | | | | |
| ISOTROPIC | Matusik et al. [4] | ✓ | ✓ | ✗ | ✓ | Provides realistic appearance and meaningful parameterisation |
| ISOTROPIC | Matusik et al. [91] | ✓ | ✗ | ✗ | ✓ | Reduces number of samples to acquire and represent BRDF |
| ISOTROPIC | Romeiro et al. [92] | ✓ | ✗ | ✓ | ✓ | Bivariate representation, allows to capture off-specular and retro-reflections. |
| ANISOTROPIC | Kautz and McCool [93] | ✗ | ✗ | ✗ | ✓ | SVD or ND based based decomposition for BRDFs; approximation based on textures, used to store directions. |
| ANISOTROPIC | McCool and Ahmad [94] | ✓ | ✓ | ✗ | ✓ | Based on logarithmic homomorphism. Simple parameterisation and limited storage cost (2 textures). |
| ANISOTROPIC | Lawrence et al. [41] | ✗ | ✗ | ✗ | ✓ | Provides accurate results and can be also used for BTFs; compact representation |
| ANISOTROPIC | Lawrence et al. [95] | ✓ | ✓ | ✓ | ✗ | Suited for interactive rendering/editing |
| ANISOTROPIC | Ozturk et al. [56] | ✓ | ✗ | ✗ | ✗ | Computationally efficient linear model for approximating BRDFs. |
| ANISOTROPIC | Bilgili et al. [97] | ✗ | ✗ | ✓ | ✓ | Recursive application of the Tucker decomposition on the error term |
| ANISOTROPIC | Pacanowski et al. [99] | ✓ | ✗ | ✓ | ✓ | Projects measured BRDFs on a 2D space and approximates them with Rational Functions; small footprint. |
| ANISOTROPIC | Ward et al. [101] | ✓ | ✗ | ✗ | ✓ | Tensor tree representation for measured BSDF data. |

Table 2.1: Some of the main properties of the BRDF models described in Section 2.3. Models are grouped by category; within each isotropic/ anisotropic sub-category models are sorted by year.

tent and intuitive material representation to represent materials efficiently, would be beneficial for Computer Graphics. Although attempts to generalise reflectance models have been made by researchers, there is still no up to date universal material representation model that can fulfill such criteria and make it possible to standardise material representation.

In the following Chapter we describe the state of the art in material appearance acquisition. The data produce by the wide range of available setups is often used as ground truth to derive new BRDF models or directly for rendering.

# Chapter 3

# Reflectance Acquisition set ups

Almost every existing BRDF model has been derived from measured data. This is particularly true in case of data driven models (Section 2.6), but also many Phenomenological and Physically-based models have been developed starting from reflectance data measured either from expensive and data intensive gonioreflectometers or more recent technologies, engineer in order to save on costs and acquisition time. In this Chapter we describe the taxonomy of the acquisition setups, defined in terms of the most important component which also roughly defines benefits and limitations.

## 3.1 Introduction to acquisition set ups

Measuring or calculating how a surface interacts with light is a time consuming and expensive procedure, which generates a vast amount of data, but it is important for realistic appearance of a material model. BRDF measurements are not only used in Computer Graphics to reproduce material reflectance, but also in many other fields such as Computer Vision (*e.g.* in object recognition applications), Aerospace (*e.g.* for optimal definition of satellite mirrors reflectance and scattering properties), Optical Engineering, Remote-Sensing (*e.g.* land cover classification, correction of view and illumination angle effects, cloud detection and atmospheric correction), Medical applications (*e.g.* diagnostics), Art (*e.g.* 3D printing), Applied Spectroscopy (*e.g.* physical condition of a surface).

The setup of a typical measurement device includes a light source to uniformly illuminate a large area of a surface and a detector to measure a small area within the illuminated region [103]. Various systems with different degrees of accuracy and costs have been constructed to measure reflectance functions, ranging from gonioreflectometers to image based measurement systems; low cost setups have also been investigated [83, 104, 105]. By dropping the assumption that a material is homogeneous and opaque, many techniques for BRDF measurement can be adapted for more complex reflectance functions (SVBRDFs, BTFs, BSSRDFs). Under certain assumptions, also setups used to acquire objects geometry through the classical photometric stereo technique [106], where the point of view is kept constant between successive images while the direction of incident illumination varies, have been successfully used to recover BRDF and SVBRDF of non-lambertian surfaces [107–113]. Some of these techniques are limited to materials with a single specular lobe [108] due to the use of optimisation algorithms to recover the parameters for the Ward isotropic BRDF [49] or require the acquisition of reference objects of known shape and with similar materials as the target [109]; to reduce the number of input pictures it has been assumed bivariate BRDFs [112] or spatial coherence of reflectance, trading spatial for angular resolution [111]. The taxonomy of the acquisition setups, detailed in the following sections, is reported in Figure 3.1.

To assess the quality of an acquisition setup it is important to derive a standardised error between the measured appearance model and the original object [114]. The distance metric $\Delta E$ is particularly suitable to measure colour differences; it is defined in the *CIE XYZ* colour space [115], a perceptually uniform space which describes the chromatic response of a standard human observer to the lighting stimulus, accounting for the incident spectral power distribution. For digital image sensors, at the heart of image based systems, the most common colour space is *sRGB*, which due to the characteristics of the Human Visual System is often prone to inaccuracies [116]. Acquired RGB values could be translated into the *CIE XYZ* and post-processed for white balancing [117], although metamerism (*i.e.* spectra that appear identical to a human observer under a certain light) would still represent a source of errors. A more robust solution would make use of a carefully characterised acquisition device (e.g. a DSLR camera), to obtain either

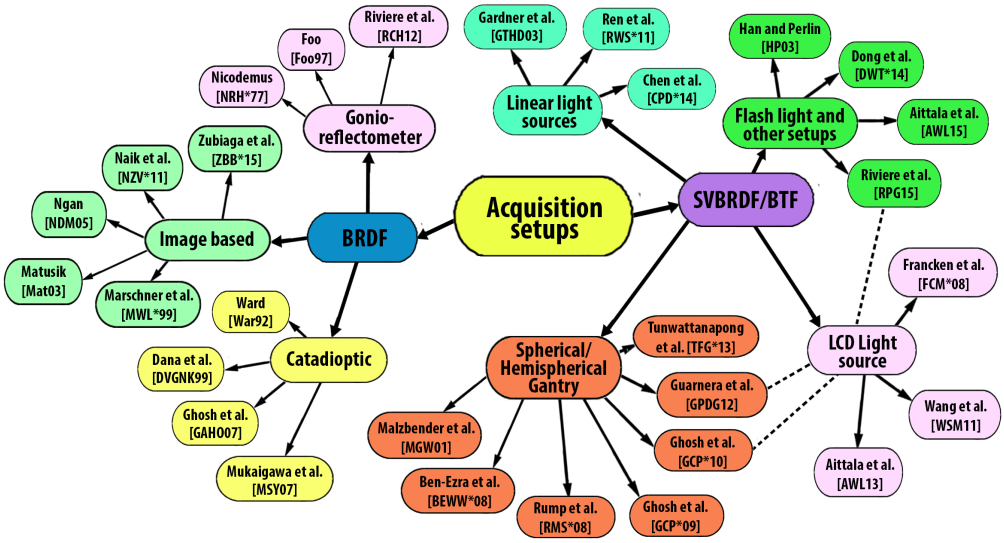


Figure 3.1: Reflectance Measurement set ups

a relative [118] or an absolute colorimetric estimation of the scene in $cd/m^2$ [119].

Numerical simulation [87, 120–123] represents, for some complex materials, a possible alternative to a measurement device. The material appearance is described by the result of the simulation of the light interaction with the surface (and sub-surface) structure. Given a geometry that can be ray-traced, Westin *et al.* in their seminal work [87] describe a method to simulate scattering hierarchically, by using the result of the simulation at a scale to generate the BRDF for a larger scale.

The measurements are always made with some uncertainty caused by technological limitations of each measurement system component, hence estimating the BRDF of a sample material with low uncertainty would be expensive.

## 3.2 Gonioreflectometers

The gonioreflectometer measures the spectral reflectance of surfaces, it covers specular and diffuse reflectance depending on the settings of the device. The construction of the device was described by Nicodemus and used in the experimental development of the reflection models by Torrance and Sparrow [59], Blinn [47],

He *et al.* [65] and many others.

A detailed setup is described by Hsia and Richmond [124]. It consists of a light source (a laser beam), a sample material placed on the sample holder mounted on a turntable which rotates around the vertical axis and a detector which captures data about reflected light from the sample. The aluminium sample holder, painted with matte black paint, is mounted on the arm attached to the turntable and placed in front of the sample detector. Two averaging spheres, with the inside part coated with barium-sulfate, are used to measure the incident light.

Foo [125] designed a three axis automated gonioreflectometer with two degrees of freedom. The measuring system consists of a light source moving around a sample, a stationary detector and a folding mirror. The system can measure the reflection at high grazing angles (up to 86 degrees) and allows high dynamic range measurements, making it considerably precise. Unfortunately, this setup can only measure isotropic BRDFs. A similar setup is described by Li *et al.* [126].

Riviere *et al.* [127] used an in-plane multispectral polarised reflectometer. The measurement setup consists of a lighting system with three linearly polarised laser sources; the polarised detection system is based on the Fresnel equation to identify polariser's axes. It allows sampling at zero lighting angles and it is fully calibrated for polarised and multispectral in-plane BRDF measurements. Polarised measurements are used to distinguish the different scattering processes in BRDF directional components. This measurement system is suggested for analysis of physical measurements of the optical surface and for laser-imaging applications. It allows users to retrieve BRDF data which prove to be numerically stable; an inversion algorithm is required for high angular measurements of BRDFs [127].

## 3.3 Image based set ups

Image-based BRDF measurement makes use of photographs of an object and requires only general-purpose equipment, thus lowering the cost of the process. The data can be measured quickly and completely through a series of photographs

taken of a surface. These photographs capture light reflected from various surface orientations. However, to measure the wavelength spectrum of the BRDF requires more time per measurement [128].

Marschner *et al.* [129] presents a rapid, complete and accurate isotropic BRDF measurement setup for a broad range of homogeneous materials, including human skin. It can achieve high resolution and accuracy over a large range of illumination and reflection directions. This setup consists of a hand-held digital camera, equipped with a standard CCD sensor with RGB colour filter array, and an industrial electronic flash light source, which suffice to measure surfaces with simple shapes, *e.g.* spherical and cylindrical which can be defined analytically; for more complex irregular shapes a 3D scanner is required in addition. The camera, characterised in terms of Optoelectronic Conversion Function (OECF) in order to know the radiance reflected to the camera and the irradiance due to the source, moves from near the light source, to measure near retro-reflections, to opposite the light source, in order to measure grazing-angle reflection. Some additional photographs are taken to measure the location and intensity of the light source, the camera pose and the sample pose. About 30 images from different positions are required to cover the three-dimensional BRDF domain. Each pixel in the images is used to derive one sample in the domain of the BRDF, thanks to the estimated relationship between the geometry of the sample and the position of the camera, light source and sample, through bundle adjustment. A typical measurement session takes up to half an hour.

A more recent development by Matusik *et al.* [4] represents full measured 3D isotropic BRDFs stored in a tabular form and interpolated linearly using a small number of basis functions to cover the entire space [56]. Matusik's measurement setup is similar to Marschner *et al.* [129], and has been used to measure 100 isotropic materials. The database, which is to date the largest and most reliable BRDF database in the field [58], is partly shown in Figure 3.2. Matusik′s data-driven method is described in Section 2.6.1. allows fitting of measured data directly in the rendering process, thus preserving measured data in high detail and resulting in a realistic material appearance. The drawback of this representation is its size; fitting data to all objects in a scene requires a huge amount of memory [130].

Ngan *et al.* [66] presented an anisotropic BRDF acquisition setup for flat and flexible samples. To deal with the anisotropy, strips of the material at different orientations obtained from flat samples are wrapped around a cylinder, which can be tilted by means of a precision motor in order to account for the missing degree of freedom with respect to a sphere. A light source rotates around the cylinder while the target is captured by a fixed camera, enabling the capture of the full 4d BRDF. For each light and target position a set of 8 pictures with different exposures is taken, to form an HDR image. The sampling density of the light and the cylinder tilting can be adjusted to increase the resolution of the measured BRDF, whereas the main limitation in the resolution is due to the limited number of material strips which can be wrapped around the cylinder.

The reflectance acquisition setup proposed by Naik *et al.* [131] exploits space-time images captured by a time-of-flight camera. Two different setups are described, both based on indirect viewing with 3-bounce scattering and making use of two known Lambertian materials, respectively the source S and the receiver R, while P is the patch to measure. In the first setup, the laser illuminates S, and the camera views R, thus measuring P indirectly. As for the second configuration, it is based on an around the corner viewing in which P is not directly visible to the camera, whereas S and R are the same surface.The light is multiplexed along different transport paths and some of them might have the same length, hence the light can arrive along multiple paths at the same point at the same time. For this reason the measurements of the material need to be decoded, by solving a sparse underdetermined system; the system is solved by recovering the parameters the Ashikhmin-Premoze model [58] (see Section 2.4), using the halfway vector parameterisation. When the multiplexing does not cause ambiguities, in order to measure the parameters of a material it is enough to analyse the streak images to find the specular peak. This setup enables to take many BRDF measurements simultaneously, but it requires an ultra-fast camera; moreover it suffers from a low signal to noise ratio due to the multiple bounces, the size of patches and the maximum sharpness of the reflectance function are limited by the hardware and the range of measurable incoming and outgoing directions is limited by the geometry of the setup.

## 3.4 Catadioptric measurement setups

Catadioptric optical systems makes use of both reflected and refracted light, in order to reduce aberrations. The resulting imaging setups are generally efficient image based BRDF acquisition devices, usually without any moving parts.

The imaging gonioreflectometer described by Ward [49] measures anisotropic surfaces by repeating measurement process under various orientations. It captures the entire hemisphere of reflected and refracted directions at the same time. Ward used a fisheye lens and half-silvered hemisphere. This device cannot measure sharp specular peaks nor take measurements at high grazing angles.

Dana *et al.*'s [132] measuring device consists of a robot arm that holds and rotates a sample, a halogen bulb with a Fresnel lens and a video camera. The light is fixed, and the camera is moved to record measurements from seven different locations. The seven location points of the measurements correspond to sample viewing points and illumination direction. For measurement of the anisotropic material the sample is rotated about the z-axis and this procedure is repeated [132]. This system was designed for use in Computer Graphics, and like Ward, includes reflection and refraction capture, however there are issues with noise within measurements and scale as the surface patches are too large to measure fine scale texture variations.

Mukaigawa *et al.* [133] built a measurement system for anisotropic BRDFs which uses a projector as the light source, placed at the focal point of an ellipsoidal mirror, a camera and a beam splitter, since the camera and the projector cannot be located at the same position. The number of acquired images depends on the sampling of the lighting direction and viewing direction, which needs to be estimated based on the accuracy required. The acquired data are then fitted to the Ward anisotropic reflection model.

Ghosh *et al.* [134, 135] describe the setup of a measurement device that does not involve any moving parts and consists of a camera focusing on a zone of reflected directions, a light source with a beam splitter, a mirrored dome and mirrored parabola. The focus of the illumination beam is on the mirrored components that the beam reflects back to its origin. This setup allows BRDF measurement over a continuous region with a specially designed orthonormal zonal basis func-

tion illumination, which results in a very rapid BRDF acquisition and in a better signal to noise ratio compared to point-sampling the incident directions [133]. The measurements are then projected into a spherical harmonics basis or fitted to an analytical reflection model.

## 3.5 Spherical and Hemispherical Gantry

Malzbender *et al.* [136] built a hemispherical device with 50 strobe light source. The camera is placed in the apex of the device, and it is used to acquire pictures of almost flat samples, placed on the floor and illuminated by a single light source at a time. The acquired data are represented by Polynomial Texture Maps (PTM), in which for each fitted texel the coefficients of the following polynomial are fitted to the data and stored as a map:

$$\begin{aligned} L(u,v;l_u,l_v) &= a_0(u,v)\,{l^2}_u + a_1(u,v)\,{l^2}_v + \\ &+ a_2(u,v)\,l_u l_v + a_3(u,v)\,l_u + a_4(u,v)\,l_v + a_5(u,v) \end{aligned} \tag{3.5.1}$$

where $L$ is the surface luminance at $(u,v)$, the local coordinates of the texture and $(l_u,l_v)$ are the projection of the normalised light vector at that coordinate. PTMs facilitate good quality rendering, in particular for diffuse samples.

A hemispherical device for anisotropic BRDF measurement was presented by Ben-Ezra *et al.* [137], in which it is demonstrated that with an accurate radiometric and geometric calibration LEDs can be used as light sources and as detectors, without needing any moving parts nor cameras; this setup allows fast acquisition times. In their implementation 84 LEDs pointing toward the centre of the hemisphere are used. During the acquisition, each LED is switched on, in turn acting as an emitter, while all others measure the reflected light from the sample. The SNR of the measurements can be increased by multiplexed illumination and the use of different colours for the LEDs allows capture of multispectral data. Since a LED cannot be used at the same time as an emitter and detector this setup cannot be used to measure retro-reflection and offers a lower resolution compared to camera-based setups.

The measurement device presented by Rump *et al.* [75] consists of a hemi-

spherical gantry with 151 cameras uniformly distributed; the cameras flashes are used as light sources and for each flash all the cameras take a picture of the subject, giving a total of 151x151 = 22,801 pictures, which can be increased by taking HDR sequences. The gantry is capable of supporting projectors in order to project structured light on the subject.

Ghosh *et al.* [138] proposed three different setups to estimate spatially varying BRDFs for both isotropic and anisotropic materials, using up to 9 polarised second order spherical gradient illumination patterns. For specular reflections, specular albedo, reflection vector and specular roughness can be directly estimated from the $0^{th}$, $1^{st}$ [139] and $2^{nd}$ order [138] statistics respectively. The first setup, suitable for roughly specular objects of any shape, is based on a LED sphere with 150 controllable lights linearly polarised, with the subject placed at the centre of the sphere. The second setup is suitable for flat objects and uses as the light source a LCD monitor, placed very close to the subject, which clearly offers a smaller coverage of incident direction but with a higher resolution than the LED sphere. The third setup makes use of a roughly specular hemisphere which reflects the light emitted by a projector on the subject placed at the centre of the hemisphere, thus allowing a dense sampling; the camera observes the subject from the apex of the hemisphere.

The analysis of the Stokes reflectance field of circularly polarised spherical illumination has been exploited by Ghosh *et al.* [140] to estimate the specular and diffuse albedo, index of refraction and specular roughness for isotropic SVBRDFs, assuming known surface orientation. Three different setups are used to demonstrate the technique, similar to the ones described in [138] but with the light sources covered with right circular polarisers. Four pictures of the subject are required to measure the Stokes field, three of them with differently oriented linear polarisers in front of the camera and one with a circular polariser.

The same framework based on the analysis of the Stokes reflectance field has been further exploited by Guarnera *et al.* [141] and it is extended to cover also unpolarised illumination, to obtain a per-pixel estimate of the surface normal from the same input data as in [140]. The proposed setup makes use of a LED sphere with 346 controllable lights unpolarised/circularly polarised; the surface normals estimation is demonstrated also with uncontrolled outdoors measurement under

overcast and hence unpolarised sky, by capturing a reference dielectric sphere in the same environment.

Tunwattanapong *et al.* [142] proposed a spinning spherical reflectance acquisition apparatus. A 1m semi-circular arc with 105 LED focused toward the centre rotates about the vertical axis at 1rpm, sweeping out continuous spherical harmonic illumination conditions. They demonstrated that 44 pictures are enough to estimate anisotropic SVBRDFs and the 3D geometry of very specular or diffuse objects. This technique further generalises the approach by Ghosh *et al.* [138], since it can be applied to higher-order spherical harmonic illumination (up to $5^{th}$ order), which allows obtaining diffuse/specular separation without relying on polarisation.

Gardner *et al.* [143] built a low cost linear light source apparatus to capture flat samples making use of a fixed camera for imaging and a structured light diode. The light source is a 50 cm long neon tube, which is translated horizontally over the surface of the subject and moved in sync with the camera acquisitions. The reflectance model used to fit the measured data is the isotropic model by Ward [49], given the camera and light source positions at each frame. The laser projects a laser stripe, which is deformed by surface variations and used in order to recover the geometry, together with two scans of the light source, in a diagonal direction. A cabin light box, with two diffused cathode tubes are used as a sample holder and to project a even diffuse white light on the surface and allows measurement of the transmitted and reflected light. Overall, the system allows recovery of the diffuse and specular colours, specular roughness, surface normals and per pixel translucency for isotropic samples.

In Ren *et al.* [144] a hand-held linear light source device, together with a BRDF chart is employed to obtain spatially varying isotropic BRDFs from a video taken with a mobile phone in LDR. The BRDF chart consists of 24 square flat tiles, with known BRDFs. The tiles are made of specular materials, except one which is a diffuse standard for camera calibration (exposure and white balance). The light source is a 40cm florescent tube, slowly moved by hand over the surface and the chart, which needs to be placed alongside. This approach requires solving a number of issues, since the camera and the light source need to be placed close to the sample and the light is moved manually. Consequently, the camera and

light position are unknown, as well as the SVBRDF of the sample. Saturated values from LDR acquisition are repaired using the values in the neighbourhood and the reflectance responses are normalised and hence aligned by a dynamic time warping algorithm. Aligned samples are then used for BRDF reconstruction.

Chen *et al.* [145] present a similar setup to Gardner *et al.* [143], scanning a linear light source over a flat sample but with the significant advantage of capturing anisotropic surface reflectance. The basic assumption is that a microfacet model can be used to model the anisotropic surface reflectance. To observe the specular reflection they modulate the illumination along the light source, by means of a transparent mask. They propose two different setups which differ in form factor and employ the same 35cm CCFL lamp and DSLR camera. The desktop form factor scanner scans a linear light source over the sample, observing the SVBRDF by means of the camera; as for the hand-held form factor scanner, the sample moves with respect to the camera and the linear light source, which instead have a fixed relative position. Finally a cylindrical lens is employed to capture in a single picture a scanline of the sample. One constant lighting pattern, together with two phase shifted sinusoidal patterns suffices to reconstruct the surface reflectance.

## 3.6 LCD Light Source

Francken *et al.* [146] make use of commodity hardware such as a LCD display and a SLR camera to recover detailed normal maps of specular objects, based on the observation that the normal of a specular pixel is the halfway vector between the light direction and the view direction. To identify the light direction among $n$ different light sources they make use of a gray code lighting patterns, by taking $O(\log_2 n)$ pictures. The accuracy of the estimated normal map depends on the number of sampled light sources.

In Aittala *et al.* [147] a low cost capture setup for SVBRDFs is presented with a similar setup as Francken *et al.* Their work relies on the design of the image formation model and uses a Fourier basis for the measurements. Isotropic BRDFs are reconstructed through Bayesian inference, since the model is analytically integrable.

The capture set up by Wang *et al.* [148], consists of a vision camera and a regu-

lar LCD, used as an area light source. It allows rapid measurement of a stationary, isotropic, glossy and bumpy surface, describing its appearance with a dual-level model, which consists of the specular and diffuse relative albedos, two surface roughness parameters and a 1D power spectrum over frequencies for visible surface bumps. Two images are required for calibration, since the LCD radiance is dependent on the viewing angle. The first one captures the surface reflection while displaying a constant gray image on the LCD and the second one is taken by attaching a mirror to the surface. To establish the pose of the surface with respect to the camera a target is placed on the surface. At the micro-scale the reflectance is characterized with the Cook-Torrance model and the distribution $D$ is assumed to be Gaussian, where the standard deviation represents the roughness; similarly at the mesoscale level roughness is approximated in terms of the standard deviation. The effect of the roughness at the microscale is assumed to be a blurring of perfect mirror reflections, whereas at the mesoscale it determines a permutation of the pixels. The surface is illuminated with a half-black, half-white image with a vertical edge, and the overall roughness is estimated by fitting a Gaussian filter that blurs the step-edge image to produce the observed one. To separate the roughness for the two different scales, all pixels are sorted by intensity and reshaped back in column-major order, thus removing the permutation induced by the mesoscale roughness; the slope of the segment obtained by averaging over the rows of the sorted image is used to estimate the microscale roughness. This approach can produce visually plausible results for highly glossy man-made indoor surfaces, including some paints, metals and plastics.

Riviere *et al.* [149] propose a mobile reflectometry solution based on a mobile device's LCD panel as extended illumination source, statically mounted at a distance of 45cm above a isotropic planar material sample, at normal incidence, in a dimly lit room. The linear polarisation of the LCD panel is exploited for diffuse/specular separation, by taking two pictures of the sample with a differently orientated plastic sheet linear polariser in front of the device camera. Albedo, surface normals and specular roughness are estimated by illuminating the sample with the same lighting patterns described in [138]. Due to the limited size of the LCD panel and the position of the front camera, this setup can only acquire $5cm \times 5cm$ area of the sample; for larger samples an appearance transfer

approach, that relies on additional measurements under natural illumination, is used.

## 3.7 Flash Illumination and other Capture Set ups

Backscattering data can be used to extract an appropriate distribution for microfacets BRDF models [58]. Based on this observation mobile devices equipped with a flash light, typically near the back camera, represent near-coaxial setups particularly useful to capture the backscatter surface reflectance to be fitted in a microfacets BRDF model [149].

Riviere *et al.* [149] mobile flash-based acquisition setup estimates the diffuse and specular albedo, surface normal and specular roughness of a planar material sample, with spatially varying isotropic surface reflectance. The back camera and flash light of a mobile device are used for a hand-held acquisition of a video in a dimly lit room, capturing data of the sample from several directions over the upper hemisphere. For reflectance calibration the diffuse grey squares of an X-Rite ColorChecker are used. The top view of the sample at normal incidence is used as a reference to register the other frames. To estimate the lighting and view directions the magnetometer/accelerometer sensors or 3D tracking can be used. The surface normal of each point is computed as the weighted average of the brightest reflection direction, the diffuse albedo is estimated as the trimmed median of the measured intensities, whereas the specular albedo is estimated from the hemispherical integral of the diffuse subtracted measurements. The specular roughness is obtained by fitting the observed backscattering profile to the GGX model [69] (see Section 2.5.1). Some blurring in the reflectance maps can be introduced by misalignments and motion blur. The limited number of lighting directions suggests the use only for rough specular materials.

Aittala *et al.* [150] mobile measurement setup for stationary materials consists of a single mobile device with on-board flash light. Given a flash-no-flash image pair of a textured material of known characteristic size, a multi-stage reconstruction pipeline allows to capture the full anisotropic SVBRDF. The input images are registered through a homography, computed from manually specified points of correspondence. The flash image provides an approximate retro-reflective mea-

surement for each pixel, that combines the effect of surface normal and BRDF, whereas the other image is used as a guide to identify points on the surface with similar local reflectance. Since there is only one observation per pixel, it is assumed that multiple points on the surface share the same reflectance properies and that can be identified under ambient lighting to be combined together. The input is organised into regular tiles approximately of the same size of the repeating texture pattern, assumed to contain a random rearrangement of the same BRDF values. A master-tile is selected for relighting and lumitexels, (*i.e.* data structures to store the geometric and photometric data of one point [151]), are obtained for it. The lumitextels are regularised using a preliminary SVBRDF fit and augmented by transferring high-frequency detail from similarly lit tiles to reduce blurring. The augmented lumitexels are used in a non-linear optimizer to fit an analytic SVBRDF model and the solution is finally reverse-propagated to the full image. This setup limits the input to the retro-reflective slice of the BRDF, hence the Fresnel effect, shadowing and masking are assumed to have typical behaviour and modelled with the BRDF model A [48] (see Section 2.4). The camera field of view represents an upper limit on the width of the specular lobes which can be observed.

The idea that the variation of the reflectance over a target forms a low dimensional manifold is exploited by Dong *et al.* [152], and describes a two-pass method to accelerate complex reflectance capture, useful for both isotropic and anisotropic flat samples. During the first phase a set of high-resolution representative single-point BRDFs is captured using a hand-held device which scans over the sample. The device consists of a pair of condenser lenses, a pinhole and a camera, aligned along the same optical axis by means of a lens tube. Six high brightness LED are used as light sources, with one light at the top and the remaining at the sides. The pinhole is placed at the focal plane of the ocular condenser lens such that the camera can image the light field of a single point on the surface, while the sample is placed at the focal plane of the field condenser lens. For each light a pair of $320 \times 240$ pixel pictures is taken, with different exposures in order to obtain a $240 \times 240$ HDR image, used for local reconstruction of BRDFs by convex linear combination in a small neighbourhood. The second phase captures a set of reflectance measurements densely over the surface, by means of a fixed

DSLR camera and a hand-held light source, about 1.5 meters away from the sample and moved in a 2D plane opposite the sample itself; a mirrored ball is used to sample the incident lighting. Up to 200 pictures are acquired and used to map the manifold derived from the fist phase over the sample surface.

The measurement device presented in [104] is based on the principle of the kaleidoscope and consists of a tapered tube whose inner walls are lined with front-surface mirrors. A single camera captures the kaleidoscopic image, in which the subimages represent the same sample seen simultaneously from many different viewpoints. The sample is illuminated by a DLP projector, which shares the optical path with the camera by means of a 45° beam splitter. The properties of the sample are measured through a sequence of pictures with different illumination images, which illuminate the sample from a known range of incoming directions due to the unique sequence of reflections from the kaleidoscopic walls. The advantages of this setting, suitable for BTFs and BSSRDFs, are the absence of moving parts which enables quick measurements and guarantee perfect registration of the measurements and the low cost; radiometric and geometric calibration need to be performed only once.

## 3.8 Conclusion

Although a broad range of solutions for material acquisition exists, there is still no straightforward and clear path to follow. Moreover, the variety of acquisition methods and the lack of a standardised material representation format tend to make difficult the use of measured data, often stored in some tabulated format not directly usable in 3D packages, that needs to be encoded and compressed in such way it is readable from a specific renderer, leaving the choice of the file format to the end user of the data.

This Chapter, along with the previous one, provide the reader with an overview of the state of the art in material representation and acquisition. In the next Chapter we describe current methods to manipulate the appearance of materials, either at a BRDF model level or from acquired data.

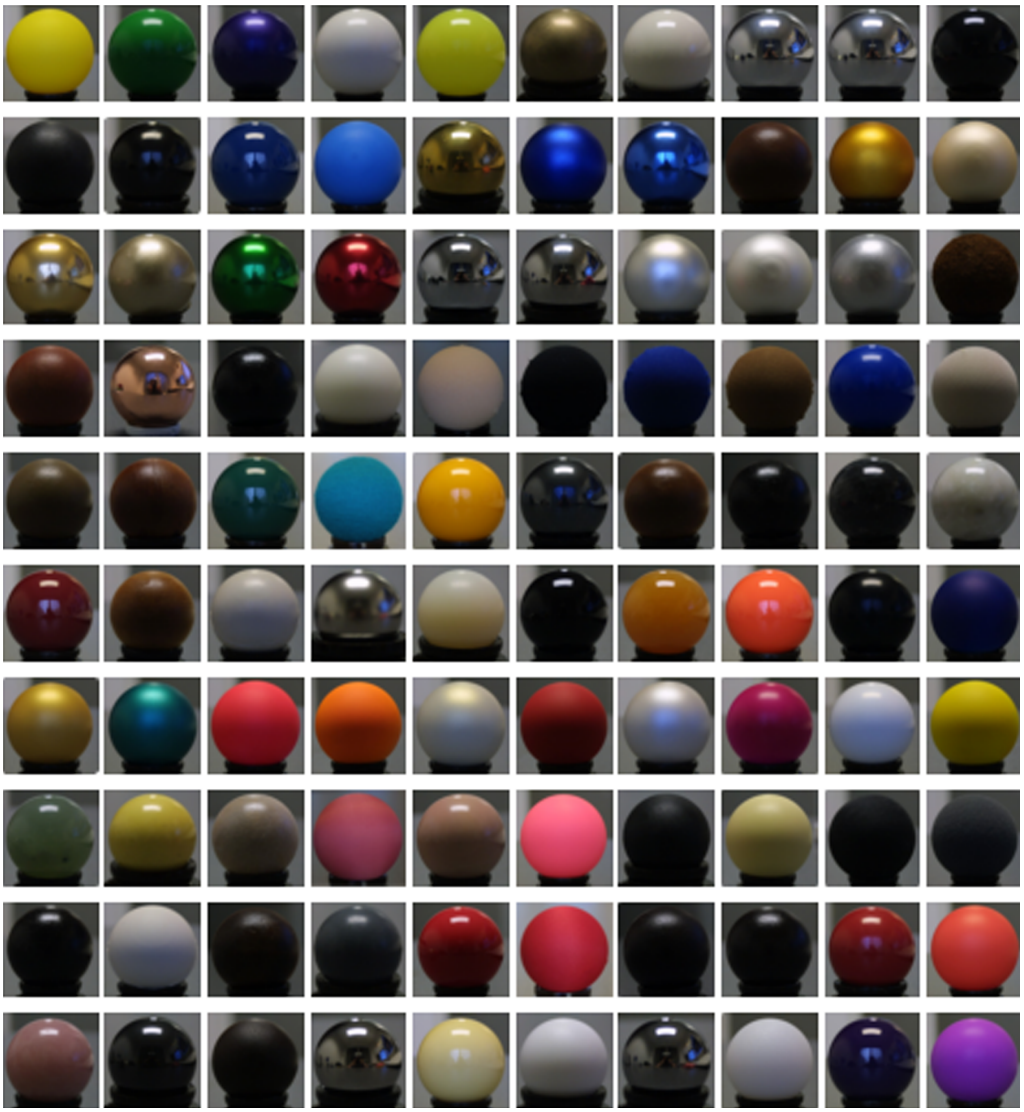

Figure 3.2: Subset of 30 materials, out of the 100 in the MERL-MIT BRDF database by Matusik *et al.* [4]. 

# Chapter 4

# Previous Work on Material Appearance Manipulation and Consistency

The previous chapters illustrates the great variety of acquisition setups used to acquire reflectance data and BRDF models to represent such data in a compact way. The broad range of material models and the complexity of the parameters requires from an artist an understanding of the underlying representation and material's micro/macrostructure; moreover current photo-realistic rendering systems use the BRDF to varying levels of accuracy, leading to a very different appearance in the final rendering. Unsurprisingly, it is very rare that the same modelling and rendering tool supports the different sets of features required by digital artists working on a marketing brochure and by a design/pre-visualisation team, the latter focusing on realistic appearance as opposed to an eye-catching one. As a result, material modelling involves a great deal of manual effort at every stage, from the acquisition to the modelling and evening in post-rendering, since the lack of support to some feature might require additional touch up, for instance with a photo editing software or with a more sophisticated appearance transfer technique.

A limitation of photo editing tools is given by the impossibility to change the material appearance at a BRDF level. Certainly it is possible to recolour an object or blur it, however changing the shape of the reflection and emulating the Fresnel

effect can only be done by manually painting pixel by pixel the appropriate areas.

This Chapter surveys current methods to manipulate the appearance of materials, either by manipulating the pixels of the rendered image or photograph (Section 4.1), by acting on the parameters of a given BRDF model before rendering an image (Section 4.2), or by manipulating the appearance of the whole image using statistics derived from similar images (Section 4.3).

## 4.1 Material Editing

Khan *et. al.* [6] leveraged some limitations of the human vision, tolerant to some physical inaccuracies, in order to develop an image-based material editing tool, which allows to automatically replace a given material with another one. It requires in input a high-dynamic range photograph of an object, together with an alpha matte to separate it from the background, to produce in output a new photograph of the same object, in which its material is replaced with an entirely new one. Some knowledge of the 3D shape of the object is required, which might not be available since the method assumes a 2D picture as input. Hence, the first step involves the estimation of a depth map, by means of an approximated shape-from-shading approach based on the pixels belonging to the object itself; the gradient of the depth map is then used to compute a surface normal $\mathbf{n}$ for each pixel. The object is removed from the image thanks to an alpha matte, either manually or automatically created; the missing pixels are inpainted by preserving the statistical properties of the remainder of the environment. A HDR environment map is then created by cutting a circle from the middle of the image, then placed in the image plane and extruded to become a hemispherical environment map. The estimated information are then used for a range of transformations ranging from applying a texture to the object to the application of an arbitrary BRDF, or even the simulation of transparency and translucency.

A MatCap (Material Capture) is the image of a sphere in orthographic projection, in which lighting and material properties are conveyed together. MatCaps represent an efficient tool for artists to design a visually plausible material appearance, removing the need to individually specify material properties and lighting. Once a MatCap is created, the appearance can be transferred to a 3D object with

a different shape, by using the surface normals as a way to map corresponding points over the surface. Thanks to these factors, MatCaps have become a successful tool, included in many physically-based renderers. However, the aforementioned benefits in the prototyping phase comes at the cost of a not easy manipulation of the material, to the extent that it might be required to recreate a new MatCap every time the material properties or the lighting need to be modified. This problem has been addressed by Zubiaga *et al.* [153], exploiting the idea that a material acts as a filter in the image, described in a previous work [154]. The input MatCap is decomposed into high and low frequency components using an heuristic strategy, based on a gray-scale morphological opening unwarped into a spherical representation, making use of the filter parameters interactively estimated as described in [154]. Overall this technique allows dynamic appearance manipulation of lighting and material, thus overcoming the typical limitation of MatCaps, able to describe only static appearance. Different MatCaps can be used on different object parts by giving in input a map of the materials IDs.

Serrano *et al.* [155] presented a new parameter space for controlling material appearance. Their BRDF editing enables an intuitive control over material and allows to create new plausible material appearance. In order to build the parameters space they extended the MERL dataset (Figure 3.2), from the initial 94 out 100 samples which are homogenous and isotropic, to 400 materials, synthesized by computing the convex hull of MERL BRDFs projected in a five-dimensional PCA. In order to uniformly cover the space with novel samples, a new BRDF is generated with a combination of the three nearest original BRDFs, weighted by their distance to the new sample. A first user study allowed to derive a set of attributes suitable to describe the appearance, which resulted in a mixture of high-level and mid-level attributes such as fabric like, soft, rough, glossy, etc. A second large scale study, performed over Amazon Mechanical Turk with 400 participants, allowed to obtain 56000 ratings (400 rendered BRDFs $\times$ 10 responses per BRDF $\times$ 14 questions per BRDF); the ratings have been used to derive one functional mapping per attribute, from the perceptual ratings to a low-dimensional representation of the BRDFs, given by the first five PCA components. Given that no a priori knowledge of the shape of the functionals, a Radial Basis Function (RBF) network is used to derive the mapping from the perceptual attributes to the

underlying PCA based representation of BRDFs. The output of the RBF is then delivered through an intuitive material appearance control graphical interface for novice user, allowing to adjust the appearance by means of a set of sliders labelled with the high-level and mid-level attributes derived from the first user study. Despite of the impressive results, a few issues can be highlighted and serve as a ground for future work:

- The MERL dataset used in Serrano [155] is well known for containing errors, in particular towards grazing angles;
- the different level of abstraction of the attributes, and their cultural dependency, might adversely affect some users;
- the functional mapping and the non clear relation between some pairs of attributes, leads to some slider (linked to the PCA-basis) conceptually difficult to understand, for instance from "fabric-like" to "metal-like", since it is not clear what kind material of materials can lie in the space spanned by these two extremes.

## 4.2 BRDF Parameters Remapping

The problem of material interchange in commercial renderers has been addressed by Sztrajman *et al.* [156]. They suggested two image-based strategies for matching the appearance of a BRDF model from a renderer (called source) to another one (called target) and remapping the parameters. This recent approach is the closest to the work we present in Chapter 5, however our research has been carried out independently and the elements in common are purely due to the fact that both algorithms address the same problem; we will highlight the differences between our approaches in Chapter 5. The framework is based on a nonlinear optimisation performed by means of the Trusted Region Reflective method, which measures the difference between source and target renderings in image space using a $L_2$ metric; starting from an initial guess of the parameters for the target model, the optimisation tries to find a set of parameters which minimises the $L_2$ metric, generating new rendering with the target model at each step. The two strategies are

called two-stage remapping, where specular and diffuse terms are remapped independently, and three-stage remapping, which builds upon the first one and makes use of its output as starting point for a third stage, in which specular and diffuse are coupled together. The idea behind the three-stage remapping is that conceptually it is better suited for layered materials such as vanished surfaces, since the assumption of independency between diffuse and specular terms, at the core of the two-stage remapping, might not hold in these cases. However their experiments show that the two-stage method is more reliable, since coupling diffuse and specular causes instabilities.

## 4.3 Appearance Transfer from Similar Images

In this Section we describe previous work on the problem of a global modification of an image. In production pipelines, very often the need of modifying the lighting or parts of the scene arises at an advanced stage of the lighting simulation (rendering), adding or removing objects and lighting effects according to the artist judgment of the preview. Unfortunately lighting simulation can take a long time to converge, and starting it all over can be impractical, since this would waste all the previous computations. For this reason, techniques which modifies the whole scene starting from a partial or complete rendering, or even a photograph, can be of great help.

Shih *et al* [157] addressed the problem of synthesizing a plausible image at a different time of day from the one in the input image, a challenging task since it often requires a huge amount of modifications, for instance when the input image depicts a daytime scene and the target time is the night. In their paper they describe a data-driven approach, which employs at its core a database of time-lapse videos of different scenes, aimed to provide information about the variations in color appearance, as a function of the time of day. The steps of the algorithm can be summarised as follows:

- Create database of time lapse videos;
- Given the input picture, retrieve a video of a scene with similar characteristics in terms of content;

- Select from the matched video the frame with the most similar histogram to the input photo, likely to represent the same time of the day;
- Compute dense local correspondence between the input and the frame from the time-lapse sequence, in order to warp the latter to the input;
- Select from the video the frame corresponding to the desired time of the day and warp it to the input;
- Hallucinating the scene colour variation over time by applying an affine colour transfer which respects the scene semantic.

This technique is able to create a plausible-looking image at a different time of day, although it might not physically correct due to shadows and highlights.

Günther *et al.* [158] describes a method which makes use of the coherence between frames before and after a scene modification to allows adding and removal of objects and light once the lighting simulation for the scene has already started, reusing much of the previous computations. This method adds a progressively computed difference image to previous frames and estimates the radiance in the modified new frame, using stochastic photon mapping as a base for the light transport algorithm. The renderings of the scenes A (pre modifications) and B (modified) can be considered as the two estimates of the pre-editing radiance and the post-editing radiance, from which the progressive difference is propagated across new intermediate frames. The sum of the pre-editing radiance and the difference image approximates the radiance in the new scene.

## 4.4 Conclusion

Rarely the same modelling and rendering tool supports the wide sets of features required by digital artists to edit material appearance. The lack of support to some feature might require additional touch up, for instance with a photo editing software or with more sophisticated appearance transfer techniques.

In this chapter we described different image based appearance manipulation methods, ranging from editing material appearance at a BRDF level, relying on

limitations of the human visual system, to appearance transfer working on the whole content of the image.

In order to achieve consistent material appearance BRDF remapping techniques represent a viable solution, by finding a mapping between the parameter space of two different BRDF models. In the next chapter we describe our automatic BRDF remapping method, which accounts for human perception and exploits a genetic algorithm optimisation driven by a computational metric in image space.

# Chapter 5

# BRDF parameters remapping

## 5.1 Introduction

Within an industry, design and development of a project are often based on virtual models. A model evolves through collaboration among several departments (for example design, marketing and development departments) and within each department digital artists are often in control of the final appearance of a model. Artists are provided with a set of materials, either in physical form or already acquired somewhere else and sent as digital data in some arbitrary file format. If a material is provided in physical form it could be left up to an artist the decision of how to measure material sample (*e.g.* scan, photograph, etc.), which can also influence the final appearance of a material.

Let us assume an artist digitalised a rubber ball with a logo on it; to represent it, a suitable rubber model might be selected, maybe mixed with a paint model. Since different rubber models might look different, the artist could try to adjust the parameters of the prebuilt model, in order to try to make it look like the physical rubber sample. Finally, the artist might send the resulting model to another department for approval, where unfortunately the final result may appear very different due to different shader and material model implementation.

Given the same 3D geometry model and a set of materials artists in different departments manually tweak materials according to their own visual perception, trying to match the original appearance provided the latter is available. The

pipeline is manual, labour expensive, time consuming and often not reusable due to 3D software updates, data exchange, and rendering with different shaders.

When a digital material sample is exchanged the overall appearance can also be affected by additional factors such as file compression and different shader implementation of a material model, since different departments might employ different 3D applications. In particular, different shaders include a variety of material presets such as plastic, rubber, paint, etc. The actual BRDF model implemented behind any of these material presets tends to be unknown, especially for commercial applications. Therefore, whenever an artist tries to match a material in different 3D application, there is no actual reference point in terms of what parameters to use, but a pure guess. Different shaders also focus on a specific subset of adjustable parameters for a given material, in order to address the actual implementation of the material model itself, which might limit reusability of a material even within the same 3D application, due to changes among shader updates.

To sum up, the 3D modelling workflow usually involves a range of modelling and rendering tools and the exchange of data between these tools is rather limited. Given a material represented with a specific model (source), not available in a different rendering tool, an artist is left with the only option to manually match the appearance using the model available in the rendering tool in use (target), *i.e.* manually finding a new set of parameters to deliver a final appearance as close as possible to the original one. This process is not a straightforward but rather time consuming and error prone.

In this chapter an automatic solution for the described scenario is given. The hypothesis of this scenario, which can be defined as "minimum-ignorance", are summarised as follows:

(i) Knowledge of the source material model used for the reference rendering;

(ii) Knowledge of the material model parameters used for the reference rendering (*e.g* roughness, index of refraction, etc);

(iii) No access to the actual implementations of the source nor the target material models.

In order to map the parameters from a source BRDF model to a target one, if we had access to the implementation of both of them (*i.e.* their source code) we could somehow reverse engineer them and provide a conversion function. However, such a solution would require programming and reverse-engineering skills which are typically beyond a typical 3D artist curriculum. Moreover, in the "minimum-ignorance" scenario, there is no access to the implementation of the source and target BRDF models (*iii*).

In this context, an artist belonging to a "department *B*" (where a "renderer *B*" is in use) could be provided by "department *A*" with a set of renderings, of a sphere for instance, rendered with the "renderer *A*" in use over there, under known lighting and model parameters. The sphere, modelled with a given BRDF model available in renderer *A* (source) but not in renderer *B* (or available even in *B* but possibly implemented in a different way), can be passed by department *B* as input to a "black-box" which can access "renderer *B*" (where the target BRDF model is implemented) and output a set of parameters for the latter, to be used in order to provide the same appearance of the source sphere.

In this Chapter, we propose a novel, automatic, image-based solution to find the best set of parameters for the target material model. As input, our method uses a source BRDF model and a few High Dynamic Range (HDR) images of a rendered sphere. Using a Genetic Algorithm (GA) based optimisation, we find a set of parameters in the target BRDF model, such that an object rendered using these parameters visually matches the appearance of the source. This GA is driven by a computational metric of similarity, defined in image-space, which compares renderings of a reference scene under specially designed incident lighting.

The main components of our proposed solution are: a reference scene, with known geometry and lighting, a set of renderings (using known parameters) of the reference scene with the source material model, a GA optimisation that can access the target renderer and material model, and a fitness function driven by a computational metric, accounting for the visual differences between the source and the target model renderings. The efficacy of our approach is evaluated through numerical validation, user studies and psychometric scaling experiments. We demonstrate that renderings of our target models are visually indistinguishable from renderings using the source BRDF model.

The chapter organized as follows: Section 5.2 provides the background on GAs, Section 5.3 describes in details the proposed framework, followed by Sectionr̃efsec:experimentsRemapping, which describes the experimental setup; Sections 5.5 and 5.6 respectively report our objective and subjective validation, followed by our conclusions 5.8. In the next chapter a different real-world scenario will be addressed, defined as "maximum-ignorance", where assumption (*ii*) does not hold.

## 5.2 Genetic Algorithms

In this section we provide a short introduction to genetic algorithm, a class of stochastic search strategies modeled after evolutionary mechanisms. Genetic algorithms offer a suitable strategy to optimise both constrained and unconstrained non-linear systems with a large number of variables, in particular when it is not possible to accurately model the interaction among them and incorporate such information into an analytical cost function [159]. Hence, genetic algorithms can be used when no information is available about the gradient of the function, which does not need to be continuous nor differentiable.

Given a target function to be optimized (fitness function) some randomly selected points in the definition domain constitute the initial population. At each step of the optimisation, the genetic algorithm accounts for the fitness function and selects individuals from the current population to be parents; from these points the next generation of children is produced. The overall effect is that, over successive generations, the population gets closer ("evolves") to an optimal solution. In analogy with the concept of evolution, at each step the following set of rules is used to create the next generation:

- one or more selection rules to identify the "parents" which contribute to the next generation;
- one or more crossover rules which establish how a new children is obtained from two parents;
- one or more mutation rules, to randomly apply changes to individual par-

ents.

There are several differences between a Genetic Algorithm (GA) and a Derivative-Based (D-B) optimization algorithm:

(1) At each step a GA generates a population of points, whereas a D-B approach generates only a single point;

(2) The best point in the population generated by a GA approaches an optimal solution, whereas the single point generated by a D-B method approaches an optimal solution;

(3) For a GA, randomisation is an important component in the selection of the next population. A D-B is strictly deterministic;

(4) A GA can deal with discontinuous and non differentiable functions, as opposed to a D-B method;

(5) Some or all parameters can be restricted to be integer values when a GA is used;

(6) A GA provides good results even when the function has several local minima or maxima, whereas a D-B method could stop at a local minima.

Given that the evaluation of a BRDF model involves a high number of variables and the exact implementation is often not available, thus making impossible to accurately model the interaction among variables, GAs provides a powerful tool to deal with BRDF parameters remapping.

The downside of GAs is the computational cost, since the function is not evaluated just at a single point a time but for all the individuals of the population; moreover multiple path are considered at the same time, rather than having a single path toward a local minima. However it is straightforward to parallelise the computations, thus increasing the performance [159]. Additional optimisations are possible, as we describes in the next sections.

## 5.3 Proposed Framework

Formally, given a distance metric $M(\cdot,\cdot) \in \mathbb{R}$ in image space, a source model $S$ controlled by $m$ parameters $\{p_{S1}, p_{S2}, \ldots, p_{Sm}\} \in \mathbf{P_S}$, a target model $T$ controlled by $n$ parameters $\{p_{T1}, p_{T2}, \ldots, p_{Tn}\} \in \mathbf{P_T}$, $m,n \in \mathbb{N}$, $m \neq n$ in general, remapping the source parameters into the target ones means finding a function (BRDF Difference Probe) $f_R : \mathbf{P_S} \rightarrow \mathbf{P_T}$ such that, $\forall \{p_{S1}, p_{S2}, \ldots, p_{Sm}\} \in \mathbf{P_S}$:

$$\begin{aligned} & f_R(\{p_{S_1}, \ldots p_{S_n}\}) = \{{p^*}_{T_1}, \ldots {p^*}_{T_n}\} = \\ & = \operatorname{argmin}\{M(I_S(p_{S_1}, \ldots p_{S_n}), I_T(p_{T_1}, \ldots p_{T_n}))\} \end{aligned} \tag{5.3.1}$$

where $I_S$ and $I_T$ are renderings of the same reference scene obtained respectively with $S$ and $T$.

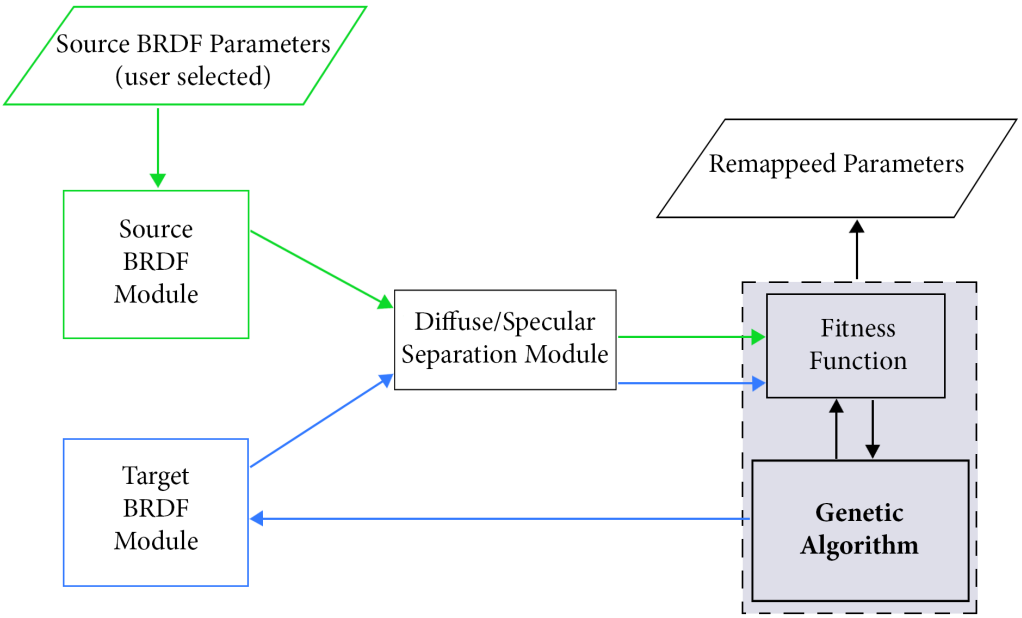


Figure 5.1: Schematic representation of the proposed automatic, unsupervised BRDF parameters remapping solution.

As reported in the Introduction to this chapter, the core of our solution is a Genetic Algorithm optimisation with access to a renderer implementing the target BRDF model (see Figure 5.1). The sphere has been selected as a reference shape for our experiment since it displays all possible surface normals. The Genetic

Algorithm starts from a random guess about the parameters of the target model and access the renderer to produce in output a set of images rendered with these parameters; driven by a suitable Fitness Function to measure the differences between the input and the output, the optimisation requests a new set of parameters to try, and another set of renderings is produced. The process is repeated until a stopping criteria is met, related to the quality of the matching between the input/output pairs. The scheme of the overall framework is reported in Figure 5.1 and the relative pseudo-code in Algorithm 1, whereas the details of each block are described in the following subsections.

### 5.3.1 Reference Scene Geometry, Lighting and Renders

We assume in input a set of renderings of a test scene containing a sphere, rendered elsewhere; the incident illumination is, in turn, a point light source and an environment map. Each of these scenes needs to be rendered several times, by changing the available parameters for the BRDF assigned to the sphere; the only requirement for the parameters selection is that their sampling spans the whole interval of meaningful values. For instance, we expect the Index of Refraction to take values in the range $]1-2.4[$ (*i.e.* from air to diamond).

The choice of the point light source, whose intensity decrease with the square of the distance and provides a smooth falloff in all direction, allows to assess how the BRDF spreads the incoming light over the surface, and it is related to parameters such as the roughness, specularity and anisotropy. On the other hand, an environment map illumination enables to describe the appearance of a sphere receiving light from all directions, such as in a natural outdoor environment, it still allows to display the effect of these parameters on the appearance of the sphere, however their analysis might prove to be more complex. The environment map includes step edges, and allows to relate the edge spread visible on the surface to the roughness of the surface itself; different colors on the environment map can allow an analysis of the surface response to different spectral bands. Hence, we employ the environment map reported in Figure 5.2, which displays sharp edges between the white, red, green, blue and black colour bands, and fine-scale color variations in the lower half, with a rich spectral content. Along with the

**Algorithm 1** Pseudo-code of the overall framework for BRDF parameters remapping. $S$ and $T$ respectively indicate the source and target BRDF models, $\{\mathbf{P}_S\}$ is the input set of $m-tuples$ of source parameters, $\{\mathbf{P^*}_T\}$ is the set of $n-tuples$ of target parameters provided in output by the algorithm. In the following sets are indicated by $\{\cdot\}$.

$\{\mathbf{P^*}_T\} \leftarrow \varnothing$
**for** each $p_{Sk} \in \{\mathbf{P}_S\}$ **do**
  $\{sourceRenderings\} \leftarrow renderImages(S, p_{Sk})$
  $\{P_{Tk,0}\} \leftarrow \varnothing$
  $p_{Tk,0} \leftarrow initRandomParameterGuess(T)$
  $\{P_{Tk,0}\} \leftarrow insert(p_{Tk,0}, \{P_{Tk,0}\})$
  $\{targetRenderings\} \leftarrow renderImages(T, p_{Tk,0})$
  $f_{k,0} \leftarrow computeFitness(\{sourceRenderings\}, \{targetRenderings\})$
  $avgFitnessPrevPopulation \leftarrow f_{k,0}$
  $\Delta_f \leftarrow avgFitnessPrevPopulation$
  $i \leftarrow 1$
  **while** $\Delta_f >$ thresh and $i <$ maximum iteration number **do**
    $\{populationFitness\} \leftarrow \varnothing$
    $\{\mathbf{P}_{Tk,i}\} \leftarrow mutatePopulation(\{\mathbf{P}_{Tk,i-1}\}, populationSize, mutationRules)$
    $f^*_k \leftarrow \infty$
    $p^*_{Tk} \leftarrow NULL$
    **for** each $p_{Tk,i} \in P_{Tk,i}$ **do**
      $\{targetRenderings\} \leftarrow renderImages(T, p_{Tk,i})$
      $f_{k,i} \leftarrow computeFitness(\{sourceRenderings\}, \{targetRenderings\})$
      $\{populationFitness\} \leftarrow insert(f_{k,i}, \{populationFitness\})$
      **if** $f_{k,i} < f^*_k$ **then**
        $f^*_k \leftarrow f_{k,i}$
        $p^*_{Tk} \leftarrow p_{Tk,i}$
      **end if**
    **end for**
    $avgFitnessPrevPopulation \leftarrow average(populationFitness)$
    $\Delta_f \leftarrow avgFitnessPrevPopulation$
  **end while**
  $\{\mathbf{P^*}_T\} \leftarrow insert(p^*_{Tk}, \{\mathbf{P^*}_T\})$
**end for**

rendered images, the information about the parameters used for it is transmitted to the genetic algorithm.

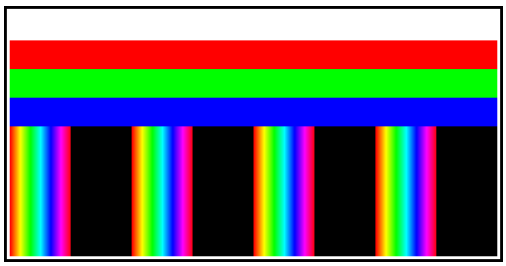

Figure 5.2: Environment map used.

A schematic representation of how the input data is produced is reported in Figure 5.3. Overall, for each set of parameters, the following images are produced:

- a diffuse only image, rendered under point light incident illumination (blue path on Figure 5.3);
- a diffuse only image, rendered under spherical lighting incident illumination (environment map, magenta path on Figure 5.3);
- a complete image (including a diffuse and a specular term), rendered under point light incident illumination (red path on Figure 5.3);
- a complete image (including a diffuse and a specular term), rendered under spherical lighting incident illumination (environment map, green path on Figure 5.3);

The same scenes are rendered upon Genetic Algorithm request, which provides in input to the renderer a set of candidate parameters for the target model.

Almost every BRDF model allows to independently specify a diffuse term and a specular term (see chapter 2), hence we exploit such characteristic to render a diffuse only image (setting all the parameters related to the specularity to zero) and a complete rendering, in which both the specular and the diffuse terms have non-zero value. This allows us to obtain a simple diffuse/specular separation by simply subtracting the diffuse image from the complete rendering, as reported in Figure 5.4; and example, rendered with the Cook-Torrance BRDF [60] (see

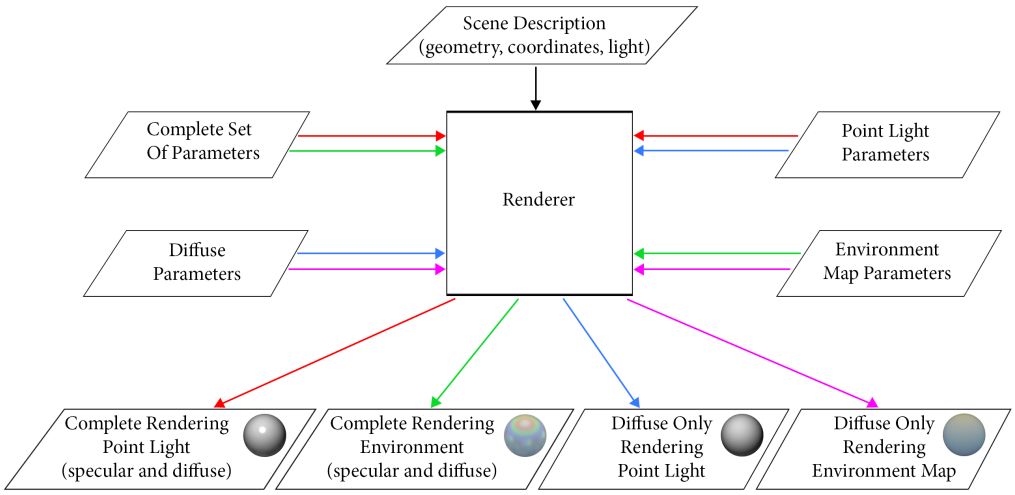


Figure 5.3: BRDF rendering module.

Section 2.5.1 ), using the Mitsuba renderer [44], is reported in Figure 5.5. This design choice has been preferred over directly rendering an image with just the specular term since it proved to be more reliable and easy to handle for dielectric materials.

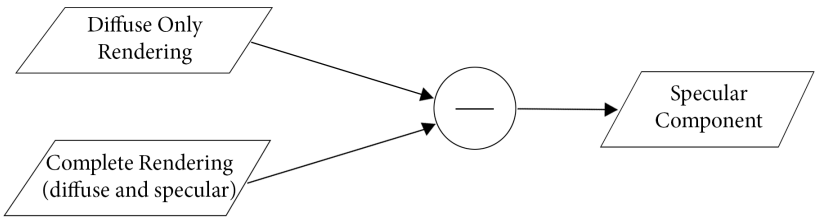


Figure 5.4: Schematic representation of the diffuse/specular separation performed in our pipeline.

### 5.3.2 Fitness Function

The fitness function has the role to establish the difference among the input and output images, for a given set of parameters, thus implementing the metric $M$ in Eq. 5.3.1. The core of the fitness function, which works with differences in the image space, is the computation of the $L_2$ metric between each couple of corresponding images (*e.g.* between the source diffuse image under point light illumination and the target diffuse image under point light illumination). The use of

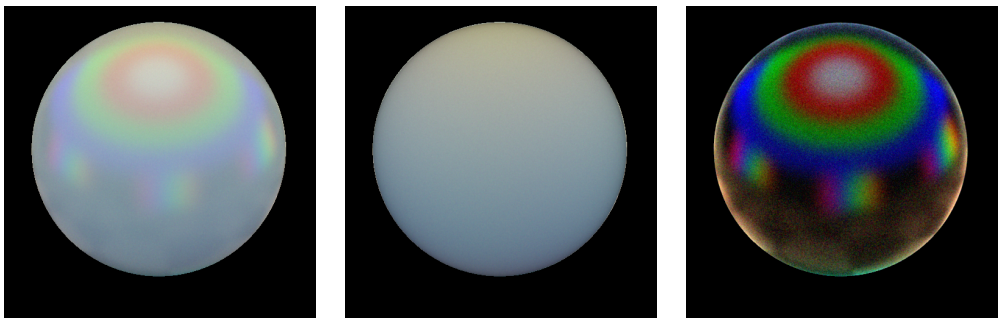

(a) (b) (c)

Figure 5.5: In this image, we show (a) a complete rendering of the sphere (with both diffuse and specular terms) under environment map illumination, (b) a rendering of the sphere with only the diffuse term, (c) the difference between (a) and (b), which gives the specular term; the difference has been amplified by a factor of 4 for a better visualization. The images have been rendered using the Cook-Torrance model, setting an index of refraction of 1.65 and a low value of surface roughness (0.1) in the Mitsuba renderer implementation.

a $L_2$ norm metric is common in Computer Graphics for fitting measured data to a BRDF model [48, 66] and for parameter remapping [156]. Another commonly used metric, proposed by Lafortune *et al.* [52], tends to over-fit near the mirror direction and is unsuitable for applications in which human perception is important. The $L_2$ itself is not be suitable to account for human perception, hence we apply several modifications to the basic scheme.

Let us use the following abbreviations:

- $D_{S_{PL}}$: diffuse source image, point light illumination;
- $D_{T_{PL}}$: diffuse target image, point light illumination;
- $D_{S_{EM}}$: diffuse source image, environment map illumination;
- $D_{T_{EM}}$: diffuse target image, environment map illumination;
- $S_{S_{PL}}$: specular source image, point light illumination;
- $S_{T_{PL}}$: specular target image, point light illumination;
- $S_{S_{EM}}$: specular source image, environment map illumination;
- $s_{T_{EM}}$: specular target image, environment map illumination;

- $\|\cdot\|_2$: $L_2$ norm.

As Figure 5.5 shows, the specular term tends to be, on the average over the all sphere, roughly one order of magnitude smaller than the diffuse term. For this reason, two weighting terms $w_D$ and $w_S$ are introduced, with $w_S > w_D$. In order to allow for more flexibility in the definition of the fitness function, a different importance can be given to each couple of corresponding terms, by introducing 4 additional scaling constants:

- $\alpha$: scaling term for diffuse image pair, point light illumination;
- $\beta$: scaling term for diffuse image pair, environment map illumination;
- $\gamma$: scaling term for specular image pair, point light illumination;
- $\delta$: scaling term for specular image pair, environment map illumination.

By putting all the terms together, the value $fv$ of the fitness function is computed with the following equation:

$$\begin{aligned} fv = w_D \left(\alpha \|D_{S_{PL}} - D_{T_{PL}}\|_2 + \beta \|D_{S_{EM}} - D_{T_{EM}}\|_2\right) + \\ + w_S \left(\gamma \|S_{S_{PL}} - S_{T_{PL}}\|_2 + \delta \|S_{S_{EM}} - S_{T_{EM}}\|_2\right). \end{aligned} \tag{5.3.2}$$

A diagram showing the computation performed by the fitness function is reported in Figure 5.6 .

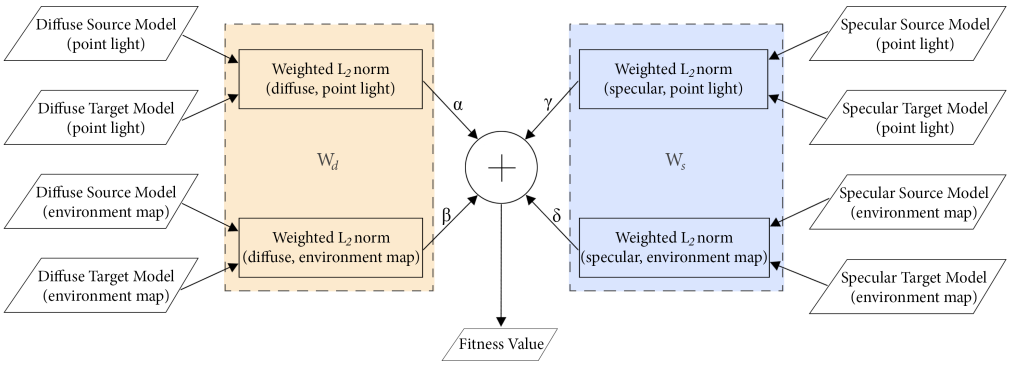


Figure 5.6: Schematic representation of the computations performed by the fitness function.

While the reflectance in dielectric materials includes both a diffuse and a specular components, and hence precisely fit the described framework, in case of conductors (like metals) there is only the specular component. However, the framework is general enough to deal with this kind of materials. In fact, the diffuse components are simply treated as null values, with the input and output set of images including only the specular pairs, the genetic algorithm requesting just the specular term from the renderer and the fitness function reducing to the following expression: $fv = w_S\left(\gamma\|S_{S_{PL}} - S_{T_{PL}}\|_2 + \delta\|S_{S_{EM}} - S_{T_{EM}}\|_2\right)$. In a similar way the case of perfect diffuse dielectrics is addressed.

### 5.3.3 Source - Target parameters fitting

Once the genetic algorithm finds a set of parameters for the target BRDF, which provide a good matching to the source ones, these are stored in a look-up table indexed by the parameters of the source BRDF. The algorithm then starts the optimisation all over, with a new set of source renderings and parameters. When all the source parameters have been remapped to the target ones, a fitting module takes care of finding a smooth, analytical relationship between the input and output parameters (see Figure 5.7), modelled with either a $2^{nd}$ or $3^{rd}$ degree polynomial; when such a relationship cannot be found, the output is a look-up table.

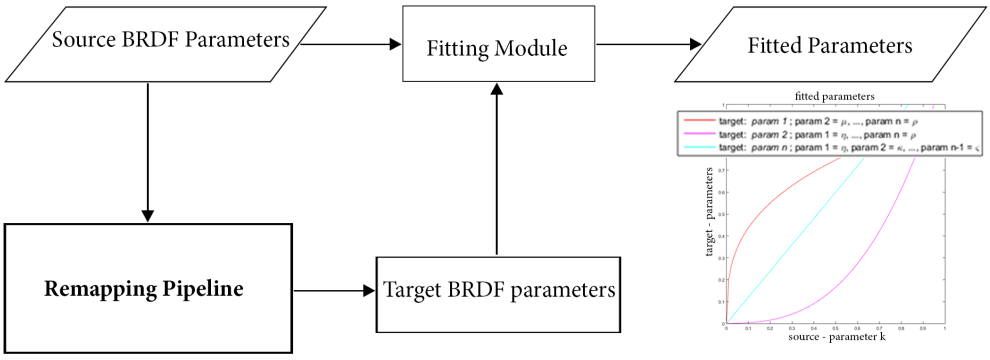


Figure 5.7: Fitting pipeline. The source images and parameters are passed in input to the remapping pipeline, one entry at a time. Once the whole list has been processed, the target parameters produced in output by the remapping pipeline are fitted to the source ones and describes by means of analytical expressions.

## 5.4 Experimental Setup

In this section we describe the implementation details of a case study involving the Mitsuba renderer, which makes available a set of different BRDFs and provides an easy interface for scripting. The genetic algorithm, fitness function and input output modules have been implemented in Matlab.

Due to the nature of the algorithm, which involves some random steps, for a given source input the fitness function can be computed several times for the same set of output parameters. In order to sped up the computations we implemented a lookup table, indexed by the values of the source and target parameters, which is initialised to null values every time a new set of source input parameters and renderings is provided to the remapping pipeline; every time a new value of the fitness function is computed, it is stored in the table for quick access.

In a similar way, we implemented a database to keep track of all the renderings produced by the target renderer (and the parameters used) upon request of the genetic algorithm: if the image is already in the database it is simply accessed and provided to the fitness function, which will search the lookup table first, looking for a precomputed values.

### 5.4.1 Dielectrics

To test the robustness of the proposed framework on dielectric materials we selected as source and target BRDF two very different models, namely the isotropic Ward BRDF, in particular the Geisler-Moroder-Dür energy conserving variant [2] and the GGX model [69] (see chapter 2). The first one is a phenomenological model, not based on the microfacets theory and with no support for the Fresnel effect, whereas the second one is a physically based model which describes a microfacets distribution and accounts for the Fresnel effect.

As explained in chapter 2, the Fresnel effect has a great impact on the appearance of a dielectric, in particular at grazing angles (see Figure 5.8); this implies that no matter what parameters are selected for the target model (GGX), the rendered images will show high errors towards the edges of the spheres when compared to the Ward rendering, since they lack the Fresnel effect. In order to

obtain a good fitting this must be taken into account, for instance by means of an extra weighting function which gives more importance to the central area of the spheres rather than the edges. In our experiments with dielectrics we implemented the weighting function by means of a mask derived with the following equation: $w_p = 1 - \left(\frac{\|p-C\|_2}{r}\right)^{\frac{1}{2}}$, where $p$ is the location coordinates of a pixel belonging to the sphere in the image lattice, $C$ are the coordinates of the center of the sphere and $r$ is the radius of the sphere; the derived weighs mask is reported in Figure 5.8(f).

In the implementation of the Ward model used for the experiments, the user can select the following parameters:

- diffuse reflectance ("$diff_{ward}$")
- specular reflectance ("$spec_{ward}$" $\in [0 - 0.45]$)
- anisotropic roughness along the tangent direction ("$alphaU_{ward}$" $\in [0 - 1]$)
- anisotropic roughness along the bitangent direction ("$alphaV_{ward}$" $\in [0 - 1]$)

under the constraint that the sum of the diffuse and specular reflectance must add up to 1 for energy conservation; the roughness along the tangent and bitangent directions are set to the same value in order to simulate an isotropic material ($alphaU_{ward} = alphaV_{ward}$). The reason for this choice is due to the lack of support for anisotropic materials in the GGX model. As for the parameters of the GGX model, the ones used in our experiments are:

- specular reflectance ("$spec_{ggx}$" $\in [0, 1]$)
- Index of Refraction of the material ("$intIOR_{ggx}$" $\in ]1 - 1.85]$)
- roughness of the surface microgeometry ("$alpha_{ggx}$" $\in [0 - 1]$)

There are other parameters available, however either they must be kept fixed for physical realism (*e.g.* specular transmittance) or are not relevant for our experiment (*e.g.* index of refraction of the means in which the sphere is immersed, air throughout all the experiments). To remap the parameters In both models the diffuse reflectance color is kept fixed to a neutral value (gray), hence we aim to find

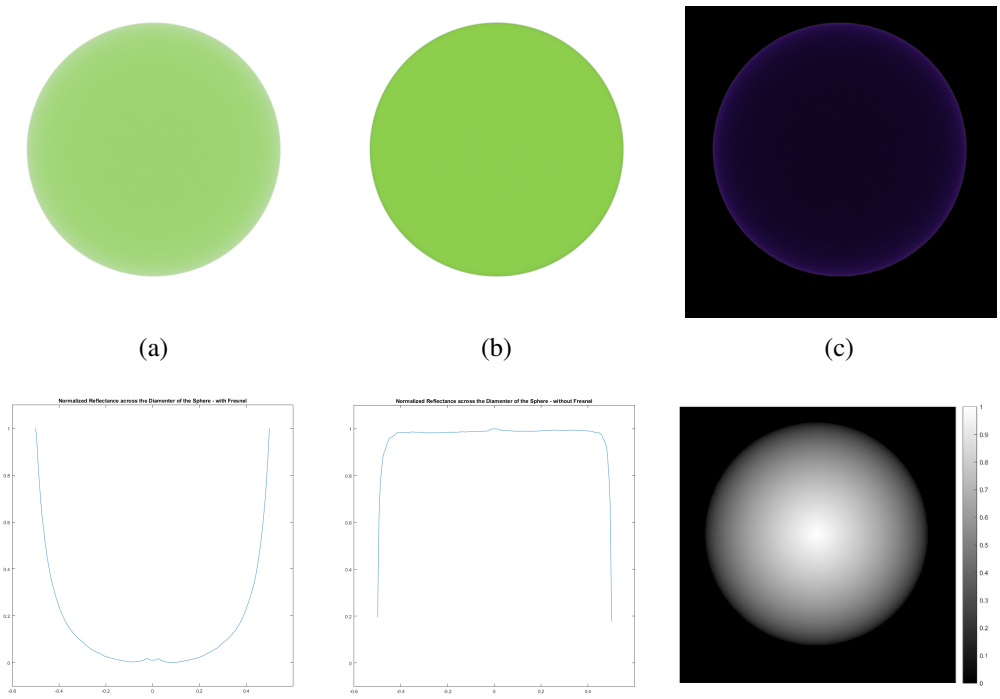


(a) (b) (c)

(d) (e) (f)

Figure 5.8: In(a) a rendering of a green polypropylene sphere under unform incident spherical lighting, index of refraction 1.49; (b) a rendering of the same sphere without accounting for the Fresnel term; (c) the difference between (a) and (b), due to the Fresnel effect and hence stronger towards the edges of the sphere; (d) normalized reflectance profile across the diameter of the sphere in (a), showing the increase of reflectance at grazing angles; (e) normalized reflectance profile across the diameter of the sphere in (b), showing that the reflectance is almost constant and decreases at grazing angles; (e) weight mask aimed to discard part of the differences between a BRDF model accounting for the Fresnel effect (*e.g.* GGX model) and one without support for it (*e.g.* Ward model).

a mapping between the following sets of parameters:
$\{spec_{ward}, alphaU_{ward} = alphaV_{ward}\} \rightarrow \{spec_{ggx}, intIOR_{ggx}, alpha_{ggx}\}$.

Finding a mapping from a pair of source parameters to a triplet of target parameters can be done in two different ways:

(A) given a pair of source parameters, allowing the genetic algorithm to explore any triplet of target parameters, searching for the best, unconstrained mapping;

(B) given a pair of source parameters, constrain the genetic algorithm to explore variation in only one dimension at a time, using knowledge from previously

remapped parameters.

While the option *(A)* is of immediate application and does not require any particular care, *(B)* involves some additional considerations and some prior knowledge of the semantic relationship among parameters must be used. It is possible to observe that:

(i) the specular reflectance in the GGX model, in order to fulfill energy conservation, can be either 0 or 1. In the Ward model this is equivalent to set the specular reflectance respectively to null or any non-null positive value;

(ii) a non-null specular reflectance in the Ward model can only be related to the index of refraction in the GGX model, which controls the Fresnel effect and hence the amount of energy reflected by means of pure specular reflection;

(iii) the isotropic roughness in the source model (Ward) can only be related to the microgeometry roughness in the target model (GGX);

(iv) the mapping between corresponding parameter is monotonically non-decreasing. If for a given source parameter $P_{ward}$, related to a target parameter $P'_{ggx}$, a mapping $\{P_{ward} = \kappa\} \rightarrow \{P'_{ggx} = \upsilon\}$ has been found, then for a value of $P_{ward} = \kappa + \delta$, $\delta > 0$ a suitable mapping can only assume greater or equal than $P'_{ggx} = \upsilon$.

To implement the above, given *(iv)* the set of source parameters must be provided in a specific order, by varying only one parameter at a time in a strictly increasing fashion. This information (along with the functional a-priori relationships *(i)-(iii)*) needs to be passed to the genetic algorithm, with the benefit of a progressive reduction of the search space and hence computational time.

Both options *(A)* and *(B)*, with the following resolution for the source and target parameters:

- Ward BRDF (source)
  - $spec_{ward} \in [0 - 0.35]$, with steps of 0.05;
  - $alphaU_{ward} = alphaV_{ward} \in [0.01 - 0.71]$, with steps of 0.05;

- GGX BRDF (target)
  - $spec_{ggx} = \{0, 1\}$;
  - $intIOR_{ggx} \in ]1.025 - 1.8]$, with steps of 0.025 for *(A)*;
    $intIOR_{ggx} \in ]1.0125 - 1.85]$, with steps of 0.0125 for *(B)*;
  - $alpha_{ggx} \in [0 - 0.5]$, with steps of 0.025 for *(A)*;
    $alpha_{ggx} \in [0 - 0.5]$, with steps of 0.0125 for *(B)*;

The higher resolution for *(B)* is due to the reduced computational cost of this configuration, as previously explained. We analysed the performances of both options *(A)* and *(B)*, by means of test scene including three spheres with different diffuse component.

### 5.4.2 Conductors

To test our framework on conductors we selected as source BRDF the phenomenological Ashikhmin-Shirley model [55], whereas the target BRDF is the physically-based Cook-Torrance model [60], with a Beckmann distribution of microfacets. The user selectable parameters of the Ashikhmin-Shirley source model used for the experiments are:

- real component of the material index of refraction (material dependent)
- imaginary component of the material index of refraction (material dependent)
- anisotropic roughness along the tangent direction ("$alphaU_{as}$" $\in [0 - 1]$)
- anisotropic roughness along the bitangent direction ("$alphaV_{as}$" $\in [0 - 1]$)

Since we focus on isotropic conductors (*i.e.* our analysis excludes, for instance, brushed metals), we set ($alphaU_{as} = alphaV_{as}$). Again, this choice is due to the lack of support to anisotropic materials in the Cook-Torrance model. As for the parameters of the Cook-Torrance model, the parameters are:

- real component of the material index of refraction (material dependent)

- imaginary component of the material index of refraction (material dependent)
- roughness of the microgeometry ("$alpha_{ct}$" $\in [0-1]$)

To perform the remapping we decided to consider real world conductors, hence the real and imaginary components of the index of refraction are determined by the material considered, (iridium to find the mapping, whereas gold, silver and copper are used to test the goodness of the mapping). To summarise, we search for a mapping between these two sets of parameters:
$\{alphaU_{as} = alphaV_{as}\} \rightarrow \{alpha_{ct}\}$.

The 1-D mapping $\{alphaU_{as} = alphaV_{as}\} \rightarrow \{alpha_{ct}\}$ we are looking for constitutes a computationally simpler case than dielectric materials; furthermore, from the latter we can also borrow observations *(iii)* and *(iv)*).

Despite of the apparent simplicity, conductors remapping allows to highlight some challenging situation, for instance when remapping a physically-based source BRDF, based on a gaussian distribution, to the target model GGX, which implements the Trowbridge-Reitz distribution [72]. We designed the remapping from Ashikhmin-Shirley to GGX as an additional experiment; the parameters available for the GGX model are the same as the Cook-Torrance model, hence we look for a 1-D mapping $\{alphaU_{as} = alphaV_{as}\} \rightarrow \{alpha_{ggx}\}$. As we discuss in the following sections, the great differences in the microfacets distributions makes the parameters remapping technique to lead to higher errors as the roughness increases.

The remapping technique has been executed with the following resolution for the source and target parameters:

- Ashikhmin-Shirley BRDF (source): $alphaU_{as} = alphaV_{as} \in [0-0.5]$, with steps of 0.01;
- Cook-Torrance BRDF (target 1): $alpha_{ct} \in [0-0.5]$, with steps of 0.001;
- GGX BRDF (target 2): $alpha_{ggx} \in [0-0.5]$, with steps of 0.001.

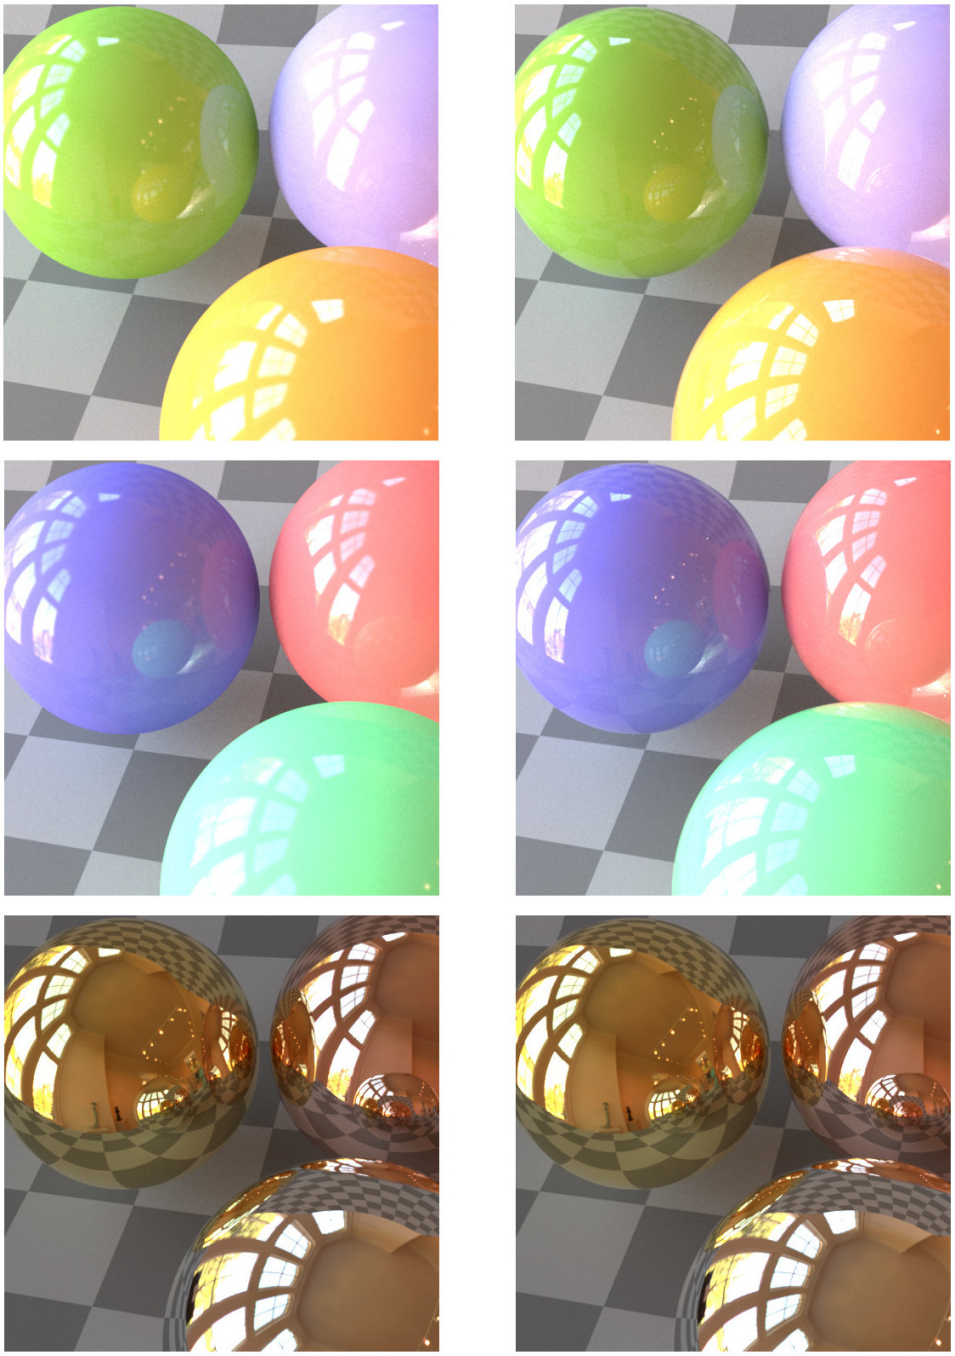

Figure 5.9: Test scenes used in our experiments. The first column shows the scene rendered with the source BRDF, the second column shows the corresponding remapped target BRDF. The first row refers to the experiment with dielectrics, in the unconstrained setting. The second rows refers to the second experiment with dielectrics, using prior knowledge of the semantic relationships among parameters. The bottom row refers to the experiment with metals.

## 5.5 Numerical Validation - Results

In this section we report a numerical validation of our remapping framework on both dielectrics and conductor, using the setups described in the previous section.

The test scene consists of three spheres with the same radius, close to each other; each sphere has a different diffuse reflectance in case of dielectrics remapping, whereas in case of conductors remapping we used three different measured metals, in particular gold, silver and copper; within the same image the BRDF parameters are the same for all the spheres, except for the aforementioned differences. The spheres are placed over a checkerboard and illuminated by an environment lighting with a rich content (windows, statues, furniture, etc.), different from the one used by the genetic algorithm to remap the parameters; an additional point light source is inserted in the scene. Some examples of the test scene are reported in Figure 5.9.

### 5.5.1 Dielectric Material

Due to the stochastic nature of both the path tracer and the sample generator, two renderings of the same material with the same parameters can slightly differ from each other. Hence, the first thing to test is whether the differences between two renderings with the same parameters are reasonably lower than the differences between two renderings obtained by changing the material parameters with the selected resolution, otherwise this could lead to instabilities in the remapping. We experimentally verified that, on the average among 10 repetitions, the difference due to the stochastic process, computed as Normalized Root Mean Square Deviation (NRMSD) is below 0.015, whereas the difference due to a change in any of the parameters by the selected step in resolution is above 0.025.

As described in Section 5.4, the first experiment with dielectrics employs an unconstrained remapping strategy from the Ward BRDF to the GGX BRDF. While for each non-null value of specularity in the Ward model we report five different values of roughness (0.01; 0.06; 0.16; 0.31; 0.71;), there is only one value of roughness for $spec_{ward} = 0$. This is due to the fact that in the latter case the reflectance reduces to the perfect diffuse lambertian model, not affected by vari-

ations in the roughness, as experimentally verified in Figure 5.10(a), where a sequence of renderings of the diffuse term with increasing roughness is reported along with the differences between each pair of two consecutive images. The NRMSD is within the range of the random changes between two renderings with the same parameters, whereas for a sphere with a non-null specular component the differences are clearly noticeable 5.10(b).

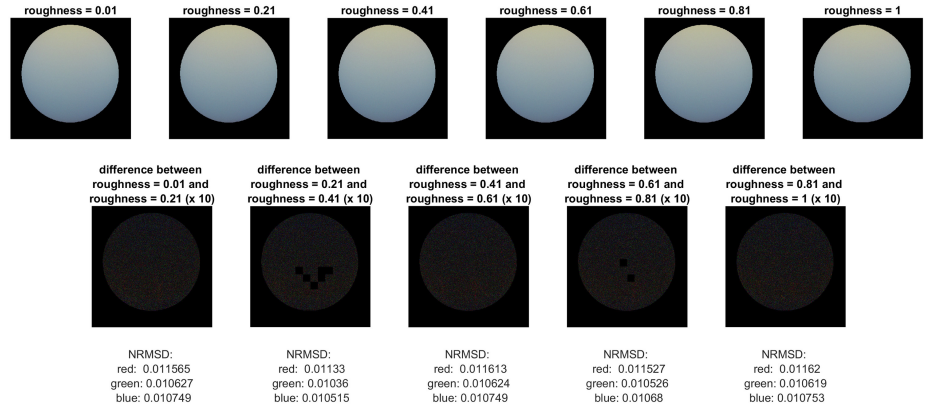


(a)

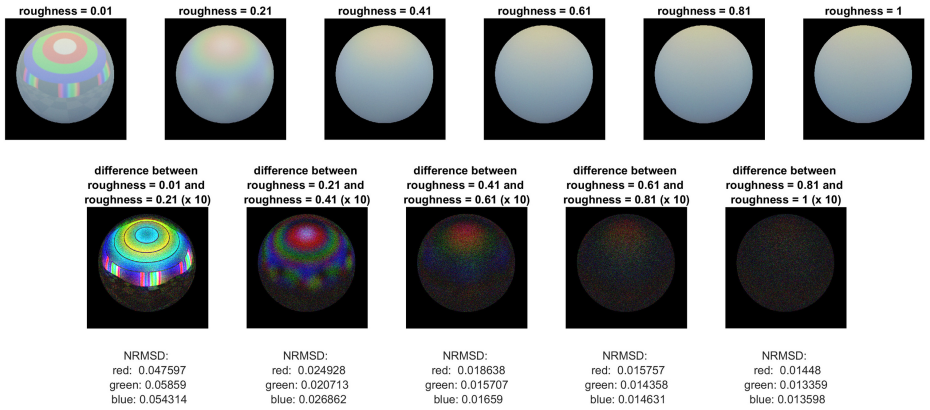


(b)

Figure 5.10: In (a), top row, we report a sequence of renderings of a sphere with $spec_{ward} = 0$ and increasing roughness (with a step of 0.2); the differences between two consecutive images in the sequence are reported in the second row, amplified by a factor of 10 to increase their visibility. In (b), we report the same kind of data for a sequence of renderings of a sphere with $spec_{ward} = 0.35$ and increasing roughness.

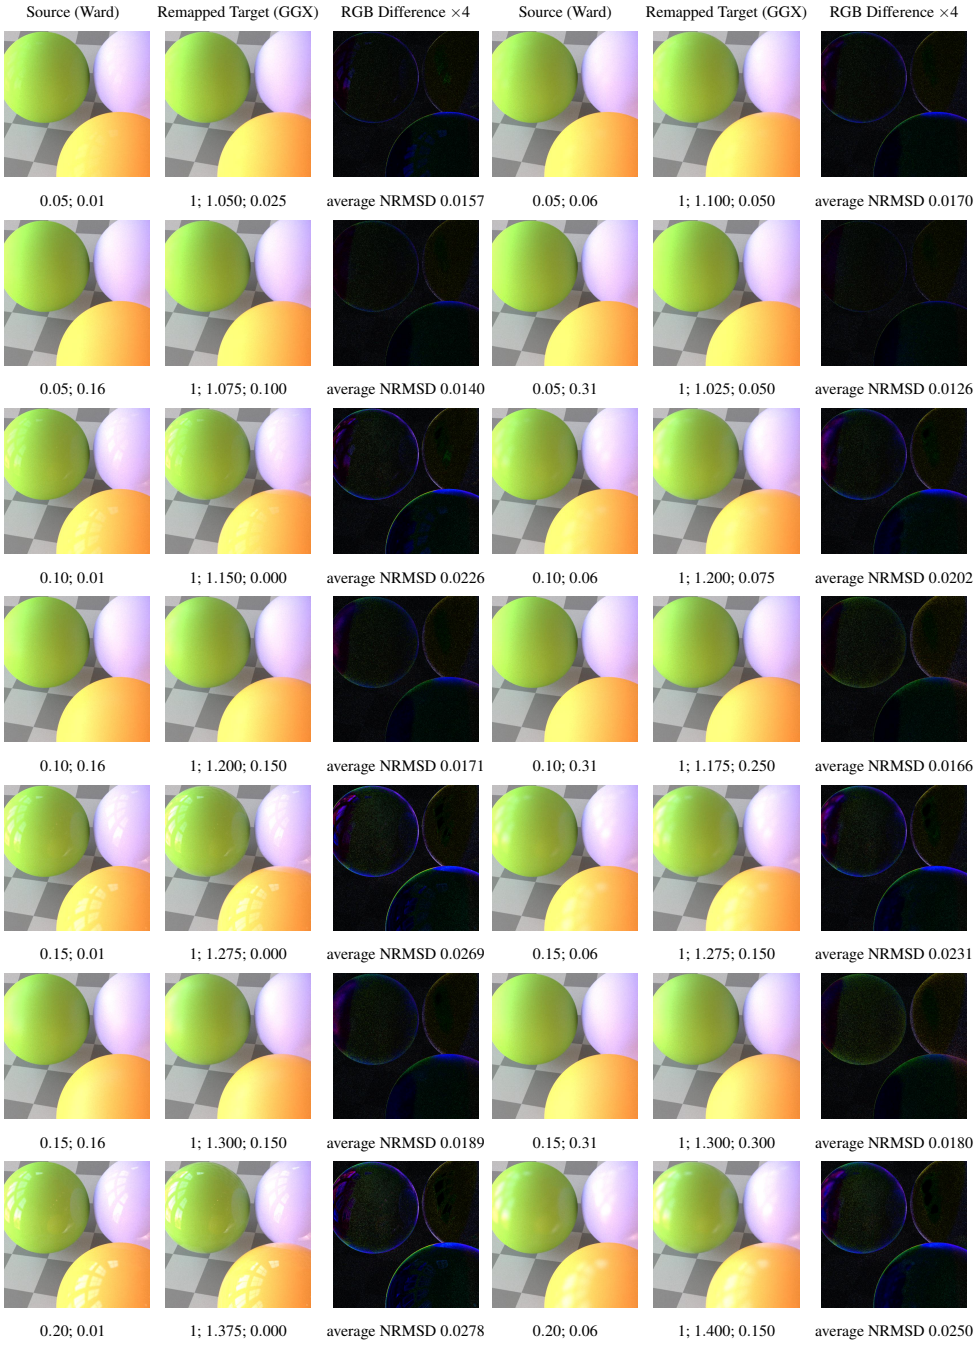


Figure 5.11: Unconstrained remapping of dielectrics. We show the source images (first and forth columns), the corresponding remapped target images (second and fifth columns) and the absolute differences in the RGB colour space, amplified by a factor of 4 for visualisation purposes (third and sixth columns). Under each of the source images we report the values of the parameters $spec_{ward}$ and $alphaU_{ward} = alphaV_{ward}$ used to render it; under the target images we report the values of the parameters $spec_{ggx}$, $intIOR_{ggx}$ and $alpha_{ggx}$ found by our framework in the unconstrained setting. Finally, under each difference image we report the average NRMSD over the RGB color channels.

In Figure 5.11 we report a visual assessment of the quality of the remapping, showing some of the source images, the corresponding remapped target images and the RGB absolute differences, along with the average NRMSD over the colour channels; table 5.1 reports the NRMSD for all the images used in the first experiment. The images generated for this experiment are also part of the user study, in order to relate the numerical error with human perception.

Table 5.1: Unconstrained remapping of dielectrics. The first column reports the $spec_{ward}$ and $alphaU_{ward} = alphaV_{ward}$ parameters of the source images, the second column reports the corresponding $spec_{ggx}$, $intIOR_{ggx}$ and $alpha_{ggx}$ parameters found by our framework. The last three columns report the Normalised Root Mean Square Deviation, per color channel.

| Source BRDF | Target BRDF | NRMSD Red | NRMSD Green | NRMSD Blue |
|---|---|---|---|---|
| (0.00 , 0.01) | (0 , 1.025 , 0.025) | 0.0115 | 0.0145 | 0.0114 |
| (0.05 , 0.01) | (1 , 1.050 , 0.025) | 0.0118 | 0.0146 | 0.0208 |
| (0.05 , 0.06) | (1 , 1.100 , 0.050) | 0.0127 | 0.0152 | 0.0232 |
| (0.05 , 0.16) | (1 , 1.075 , 0.100) | 0.0113 | 0.0144 | 0.0164 |
| (0.05 , 0.31) | (1 , 1.025 , 0.050) | 0.0105 | 0.0133 | 0.0139 |
| (0.05 , 0.71) | (1 , 1.025 , 0.025) | 0.0118 | 0.0146 | 0.0123 |
| (0.10 , 0.01) | (1 , 1.150 , 0.000) | 0.0165 | 0.0187 | 0.0326 |
| (0.10 , 0.06) | (1 , 1.200 , 0.075) | 0.0146 | 0.0168 | 0.0292 |
| (0.10 , 0.16) | (1 , 1.200 , 0.150) | 0.0134 | 0.0162 | 0.0218 |
| (0.10 , 0.31) | (1 , 1.175 , 0.250) | 0.0154 | 0.0192 | 0.0152 |
| (0.10 , 0.71) | (1 , 1.025 , 0.025) | 0.0126 | 0.0150 | 0.0154 |
| (0.15 , 0.01) | (1 , 1.275 , 0.000) | 0.0191 | 0.0216 | 0.0399 |
| (0.15 , 0.06) | (1 , 1.275 , 0.050) | 0.0163 | 0.0188 | 0.0344 |
| (0.15 , 0.16) | (1 , 1.300 , 0.150) | 0.0142 | 0.0170 | 0.0256 |
| (0.15 , 0.31) | (1 , 1.300 , 0.300) | 0.0166 | 0.0211 | 0.0165 |
| (0.15 , 0.71) | (1 , 1.225 , 0.400) | 0.0242 | 0.0289 | 0.0175 |
| (0.20 , 0.01) | (1 , 1.375 , 0.000) | 0.0196 | 0.0226 | 0.0411 |
| (0.20 , 0.06) | (1 , 1.400 , 0.050) | 0.0171 | 0.0206 | 0.0374 |
| (0.20 , 0.16) | (1 , 1.425 , 0.150) | 0.0144 | 0.0181 | 0.0285 |
| (0.20 , 0.31) | (1 , 1.400 , 0.275) | 0.0158 | 0.0203 | 0.0187 |
| (0.20 , 0.71) | (1 , 1.350 , 0.475) | 0.0287 | 0.0340 | 0.0188 |
| (0.25 , 0.01) | (1 , 1.525 , 0.000) | 0.0223 | 0.0271 | 0.0415 |
| (0.25 , 0.06) | (1 , 1.525 , 0.050) | 0.0177 | 0.0223 | 0.0371 |
| (0.25 , 0.16) | (1 , 1.550 , 0.150) | 0.0145 | 0.0192 | 0.0288 |
| (0.25 , 0.31) | (1 , 1.550 , 0.275) | 0.0157 | 0.0208 | 0.0209 |
| (0.25 , 0.71) | (1 , 1.475 , 0.500) | 0.0303 | 0.0358 | 0.0191 |
| (0.30 , 0.01) | (1 , 1.675 , 0.000) | 0.0243 | 0.0302 | 0.0404 |
| (0.30 , 0.06) | (1 , 1.700 , 0.050) | 0.0182 | 0.0244 | 0.0370 |
| (0.30 , 0.16) | (1 , 1.700 , 0.125) | 0.0156 | 0.0213 | 0.0313 |
| (0.30 , 0.31) | (1 , 1.700 , 0.275) | 0.0164 | 0.0220 | 0.0218 |
| (0.30 , 0.71) | (1 , 1.600 , 0.500) | 0.0310 | 0.0363 | 0.0197 |
| (0.35 , 0.01) | (1 , 1.800 , 0.000) | 0.0254 | 0.0309 | 0.0390 |
| (0.35 , 0.06) | (1 , 1.800 , 0.025) | 0.0245 | 0.0293 | 0.0440 |
| (0.35 , 0.16) | (1 , 1.800 , 0.125) | 0.0189 | 0.0238 | 0.0290 |
| (0.35 , 0.31) | (1 , 1.800 , 0.250) | 0.0194 | 0.0241 | 0.0220 |
| (0.35 , 0.71) | (1 , 1.750 , 0.500) | 0.0323 | 0.0374 | 0.0206 |

As for the second experiment, in which prior knowledge of the semantic relationships among parameters of the source and target models, Figure 5.12 reports a visual assessment of the quality of the remapping, in the same way as reported for 5.11. In both experiments most of the errors are localised around the edges of the spheres, as shown in the third and sixth columns of figures 5.11 and 5.12. This

is due to the lack of support to the Fresnel effect in the Ward model and cannot be fixed by a remapping algorithm, as explained in previous sections; further comparison, showing that the error is mainly localised around the silhouette and due to the Fresnel reflectance, are reported in Figure 5.13. To numerically compare the performances of the two remapping strategies, in table 5.3 we report the average NRMSD over the RGB color channels. As expected the unconstrained remapping performs slightly better than the constrained one, being the average the NRMSD 3.78% lower. The strategy which makes use of prior knowledge outperforms the unconstrained strategy in 13 images over 36 (36.1%), partly thanks to the higher resolution in the parameters space allowed by (iv).

Beyond a similar NRMSD, an important advantage of the second strategy is the possibility to derive smooth relationships between corresponding parameters, whereas the unconstrained remapping strategy forces the fitting module to simply output a look-up table, since all the target parameters are allowed to vary at the same time. This is particular noticeable for the selection of diffuse only/diffuse and specular 5.14(a) and the roughness 5.14(c) in the remapped GGX parameters, but also the remapped Index of Refraction 5.14(b) displays dependencies on both the Ward specularity and roughness. As a further consequence, it is necessary to sample the whole parameters space of the source BRDF model, since at any stage of the remapping is possible to predict how a given set of input parameters would be remapped.

On the contrary, the strategy making use of prior knowledge conveys some important generalisation characteristics, which allow to sample only a relatively small portion of the source model parameter space, as reported in Figures 5.14(e) and 5.14(f), where the genetic algorithm has sampled the source roughness up to a value of 0.45. The output of the genetic algorithm is then used by the fitting module to derive the polynomials plotted in Figure 5.15; one polynomial is enough to remap the Ward specularity into the Index of Refraction in GGX, regardless the surface roughness.

As for the roughness itself, a different polynomial remapping the source roughness into the target one can be derived for each specular level, although all the polynomials show a similar behaviour within the range $0 \leq alphaU_{ward} = alphaV_{ward} \leq 0.2$. Once a set of continuous and differentiable remapping functions has been de-

rived, they can also be inverted and used for the inverse remapping, *i.e. GGX* $\rightarrow$ *Ward* in this case.

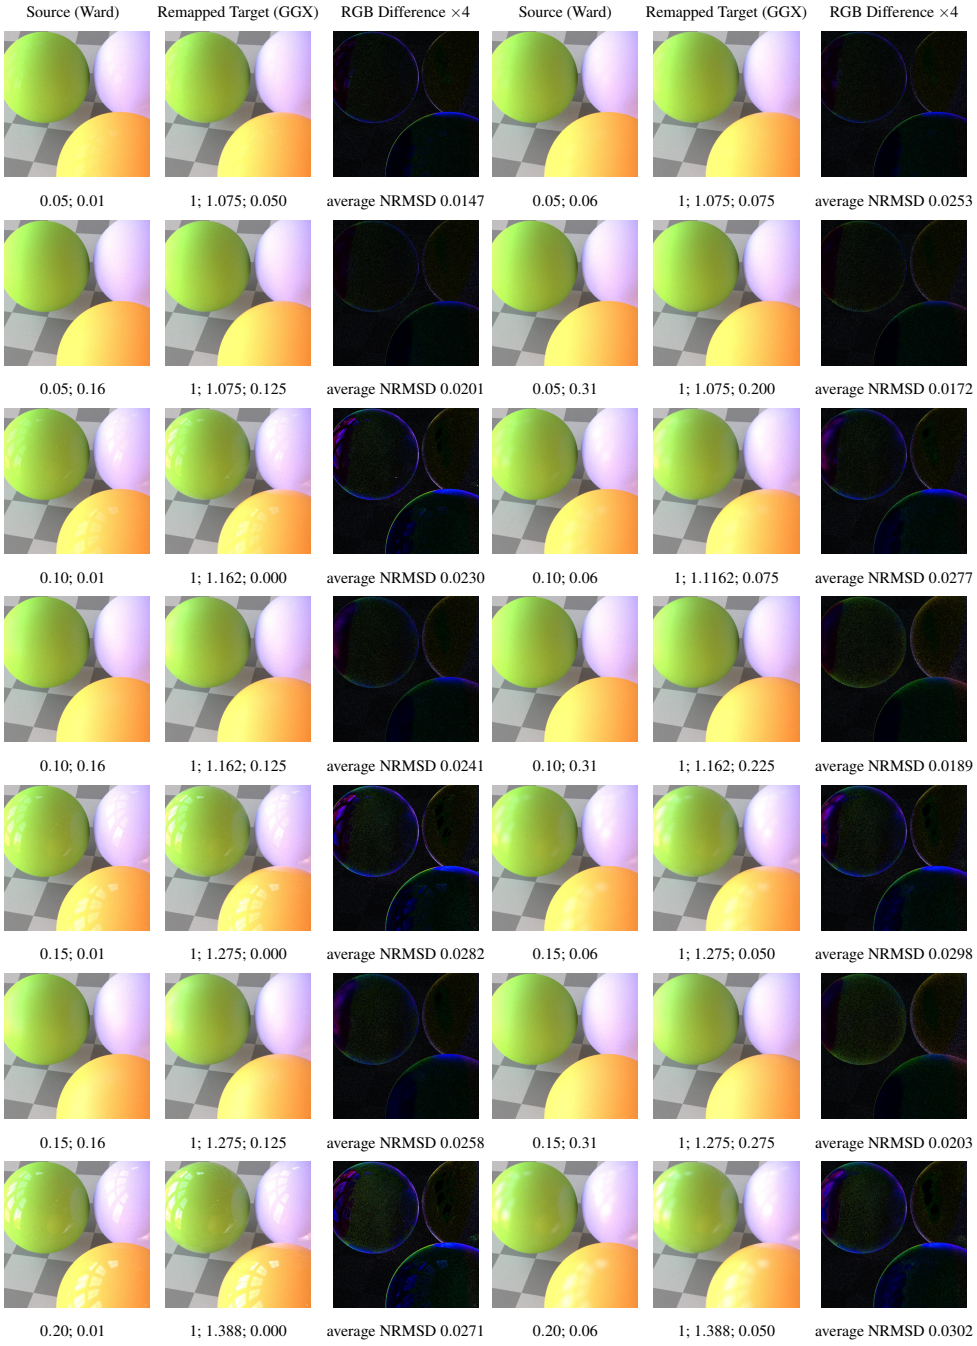


Figure 5.12: Constrained remapping of dielectrics, using prior knowledge (the same set of colours as the unconstrained case is used). We show the source images (first and forth columns), the corresponding remapped target images (second and fifth columns) and the absolute differences in the RGB colour space, amplified by a factor of 4 for visualisation purposes (third and sixth columns). Under each of the source images we report the values of the parameters $spec_{ward}$ and $alphaU_{ward} = alphaV_{ward}$ used to render it; under the target images we report the values of the parameters $spec_{ggx}$, $intIOR_{ggx}$ and $alpha_{ggx}$ found by our framework in the unconstrained setting. Finally, under each difference image we report the average NRMSD over the RGB color channels.

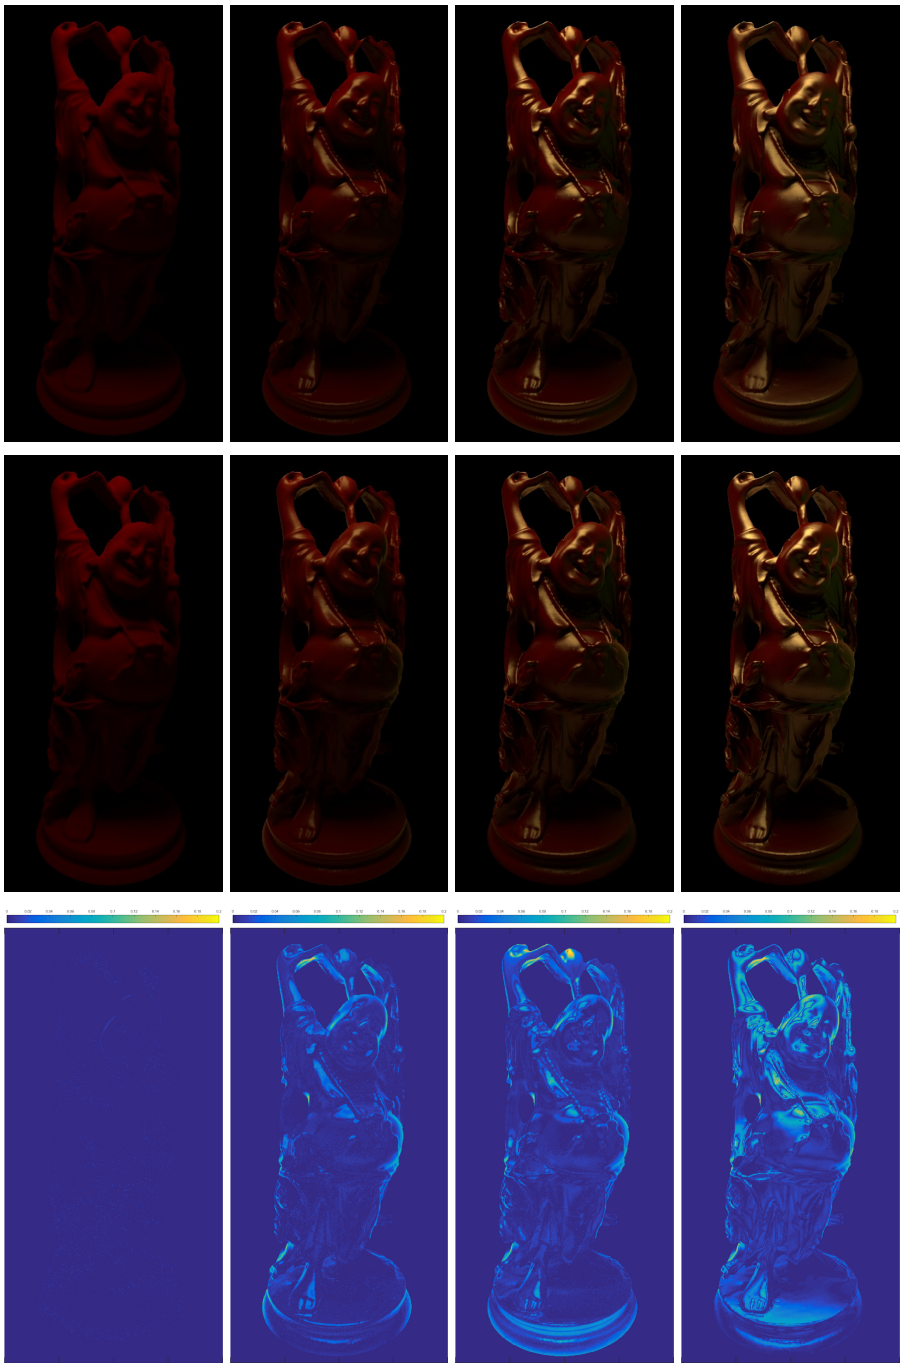


Figure 5.13: Effect of the Fresnel reflectance on the remapping error. The top row show a set of renderings in the Ward model, with increasing specular reflectance; the middle row shows the corresponding renderings with the GGX model; the bottom row shows the error, which increases as the specular reflectance in Ward increases (and the corresponding remapped IOR in GGX), and it is localised around the silhouette.

Table 5.2: Remapping of dielectrics using prior knowledge of the functional relationships among parameters in the source and target models. The first column reports the $spec_{ward}$ and $alphaU_{ward} = alphaV_{ward}$ parameters of the source images, the secondo column reports the corresponding $spec_{ggx}$, $intIOR_{ggx}$ and $alpha_{ggx}$ parameters found by our framework. The last three columns report the Normalised Root Mean Square Deviation, per color channel.

| Source BRDF | Target BRDF | NRMSD Red | NRMSD Green | NRMSD Blue |
|---|---|---|---|---|
| (0.00 , 0.01) | (0 , 1.013 , 0.025) | 0.0116 | 0.0146 | 0.0114 |
| (0.05 , 0.01) | (1 , 1.075 , 0.050) | 0.0135 | 0.0159 | 0.0245 |
| (0.05 , 0.06) | (1 , 1.075 , 0.075) | 0.0117 | 0.0146 | 0.0203 |
| (0.05 , 0.16) | (1 , 1.075 , 0.125) | 0.0113 | 0.0144 | 0.0163 |
| (0.05 , 0.31) | (1 , 1.075 , 0.200) | 0.0120 | 0.0152 | 0.0141 |
| (0.05 , 0.71) | (1 , 1.075 , 0.400) | 0.0136 | 0.0171 | 0.0135 |
| (0.10 , 0.01) | (1 , 1.162 , 0.000) | 0.0181 | 0.0198 | 0.0379 |
| (0.10 , 0.06) | (1 , 1.162 , 0.075) | 0.0146 | 0.0167 | 0.0291 |
| (0.10 , 0.16) | (1 , 1.162 , 0.125) | 0.0135 | 0.0163 | 0.0219 |
| (0.10 , 0.31) | (1 , 1.162 , 0.225) | 0.0161 | 0.0201 | 0.0156 |
| (0.10 , 0.71) | (1 , 1.162 , 0.438) | 0.0237 | 0.0286 | 0.0168 |
| (0.15 , 0.01) | (1 , 1.275 , 0.000) | 0.0193 | 0.0215 | 0.0424 |
| (0.15 , 0.06) | (1 , 1.275 , 0.050) | 0.0168 | 0.0191 | 0.0365 |
| (0.15 , 0.16) | (1 , 1.275 , 0.125) | 0.0142 | 0.0170 | 0.0256 |
| (0.15 , 0.31) | (1 , 1.275 , 0.275) | 0.0165 | 0.0210 | 0.0166 |
| (0.15 , 0.71) | (1 , 1.275 , 0.523) | 0.0301 | 0.0358 | 0.0186 |
| (0.20 , 0.01) | (1 , 1.388 , 0.000) | 0.0209 | 0.0248 | 0.0437 |
| (0.20 , 0.06) | (1 , 1.388 , 0.050) | 0.0173 | 0.0209 | 0.0392 |
| (0.20 , 0.16) | (1 , 1.388 , 0.150) | 0.0144 | 0.0181 | 0.0285 |
| (0.20 , 0.31) | (1 , 1.388 , 0.275) | 0.0162 | 0.0210 | 0.0187 |
| (0.20 , 0.71) | (1 , 1.388 , 0.527) | 0.0285 | 0.0344 | 0.0184 |
| (0.25 , 0.01) | (1 , 1.525 , 0.000) | 0.0217 | 0.0266 | 0.0424 |
| (0.25 , 0.06) | (1 , 1.525 , 0.050) | 0.0175 | 0.0225 | 0.0384 |
| (0.25 , 0.16) | (1 , 1.525 , 0.150) | 0.0145 | 0.0192 | 0.0287 |
| (0.25 , 0.31) | (1 , 1.525 , 0.275) | 0.0157 | 0.0208 | 0.0208 |
| (0.25 , 0.71) | (1 , 1.525 , 0.592) | 0.0281 | 0.0342 | 0.0188 |
| (0.30 , 0.01) | (1 , 1.663 , 0.000) | 0.0255 | 0.0320 | 0.0410 |
| (0.30 , 0.06) | (1 , 1.663 , 0.050) | 0.0182 | 0.0244 | 0.0369 |
| (0.30 , 0.16) | (1 , 1.663 , 0.125) | 0.0156 | 0.0214 | 0.0313 |
| (0.30 , 0.31) | (1 , 1.663 , 0.275) | 0.0163 | 0.0220 | 0.0217 |
| (0.30 , 0.71) | (1 , 1.663 , 0.629) | 0.0279 | 0.0345 | 0.0197 |
| (0.35 , 0.01) | (1 , 1.838 , 0.000) | 0.0254 | 0.0309 | 0.0390 |
| (0.35 , 0.06) | (1 , 1.838 , 0.025) | 0.0245 | 0.0293 | 0.0440 |
| (0.35 , 0.16) | (1 , 1.838 , 0.125) | 0.0190 | 0.0240 | 0.0292 |
| (0.35 , 0.31) | (1 , 1.838 , 0.275) | 0.0194 | 0.0241 | 0.0220 |
| (0.35 , 0.71) | (1 , 1.838 , 0.629) | 0.0306 | 0.0366 | 0.0196 |

Table 5.3: Numerical comparison between the two remapping strategies (experiments with dielectric materials): unconstrained and using prior knowledge of the functional relationships among parameters in the source and target models. The first column reports the row ID in tables 5.1 and 5.2; the second column reports the average NRMSD over the RGB channels for each of the images in 5.1 (unconstrained setting), the second column reports the average NRMSD over the RGB channels for each of the images in 5.2 (using prior knowledge).

| | Remapping strategy, average NRMSD | | |
|---|---|---|---|
| Row ID | Unconstrained | Prior Knowledge | Variation (in %) |
| 01 | 0.0125 | 0.0125 | 0.34 |
| 02 | 0.0157 | 0.0180 | 14.31 |
| 03 | 0.0170 | 0.0155 | -8.82 |
| 04 | 0.0140 | 0.0140 | -0.22 |
| 05 | 0.0126 | 0.0138 | 9.53 |
| 06 | 0.0129 | 0.0147 | 14.29 |
| 07 | 0.0226 | 0.0253 | 11.74 |
| 08 | 0.0202 | 0.0201 | -0.29 |
| 09 | 0.0171 | 0.0172 | 0.36 |
| 10 | 0.0166 | 0.0172 | 3.92 |
| 11 | 0.0144 | 0.0230 | 60.74 |
| 12 | 0.0269 | 0.0277 | 3.14 |
| 13 | 0.0231 | 0.0241 | 4.15 |
| 14 | 0.0189 | 0.0189 | 0.05 |
| 15 | 0.0180 | 0.0181 | 0.01 |
| 16 | 0.0235 | 0.0282 | 19.81 |
| 17 | 0.0278 | 0.0298 | 7.35 |
| 18 | 0.0250 | 0.0258 | 3.04 |
| 19 | 0.0203 | 0.0203 | -0.01 |
| 20 | 0.0183 | 0.0186 | 1.89 |
| 21 | 0.0272 | 0.0271 | -0.19 |
| 22 | 0.0303 | 0.0302 | -0.23 |
| 23 | 0.0257 | 0.0261 | 1.71 |
| 24 | 0.0208 | 0.0208 | -0.13 |
| 25 | 0.0191 | 0.0191 | -0.33 |
| 26 | 0.0284 | 0.0270 | -4.90 |
| 27 | 0.0316 | 0.0328 | 3.70 |
| 28 | 0.0265 | 0.0265 | -0.09 |
| 29 | 0.0227 | 0.0228 | 0.11 |
| 30 | 0.0201 | 0.0200 | -0.24 |
| 31 | 0.0290 | 0.0274 | -5.60 |
| 32 | 0.0317 | 0.0318 | 0.06 |
| 33 | 0.0326 | 0.0326 | -0.08 |
| 34 | 0.0239 | 0.0240 | 0.58 |
| 35 | 0.0218 | 0.0219 | 0.13 |
| 36 | 0.0301 | 0.0289 | -3.88 |

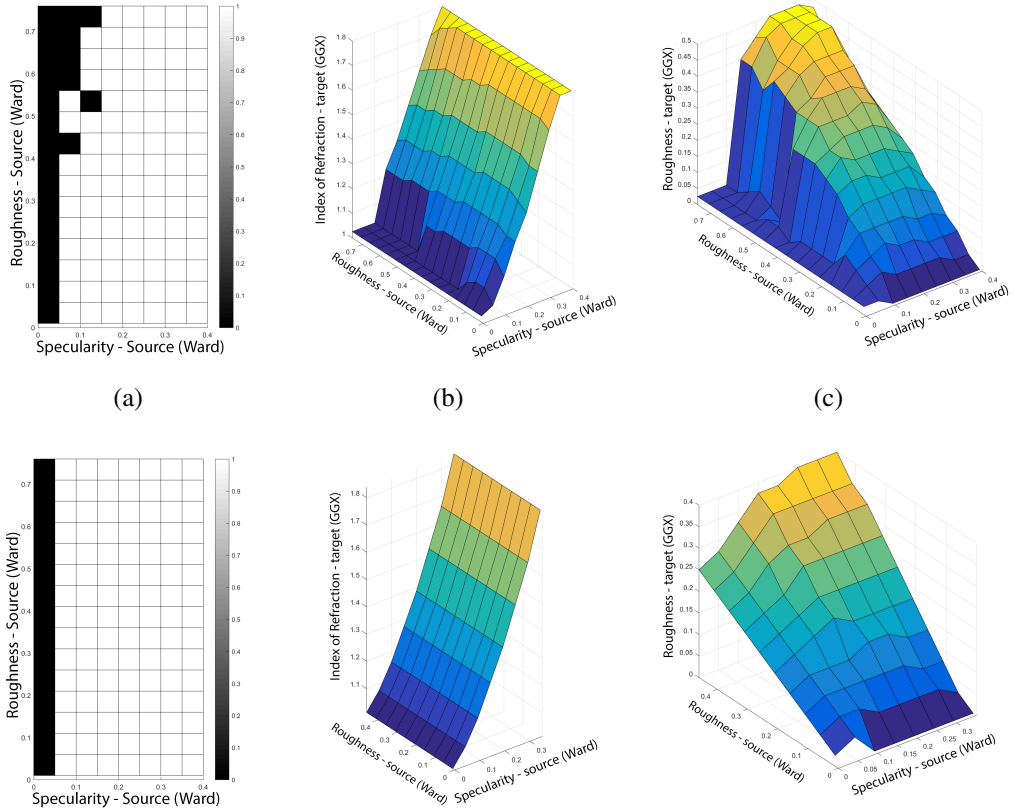


Figure 5.14: Unconstrained remapping output compared to the output of the remapping using prior knowledge. In the first row we report the output of the remapping framework in the unconstrained configuration, showing the binary choice between diffuse only (black) or diffuse and specular (white) for the remapped target parameter $spec_{ggx}$ (in (a)), the remapped target $intIOR_{ggx}$ (b) and target $alpha_{ggx}$ (c), depending on the source $spec_{ward}$ and $alphaU_{ward} = alphaV_{ward}$. The second rows reports the same information as derived by the remapping using prior knowledge, respectively in (d), (e) and (f).

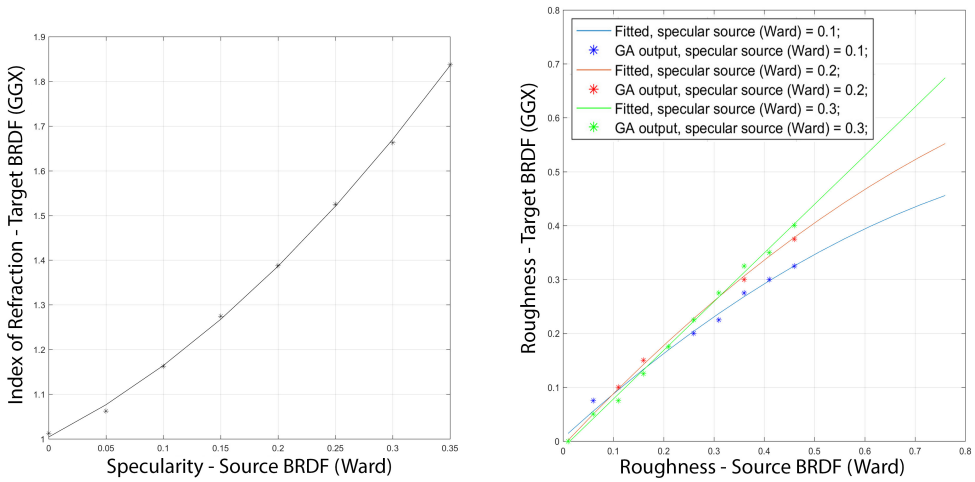


(a) (b)

Figure 5.15: Output of the fitting module, configuration using prior knowledge. In (a) the derived relation between the specularity in the Ward model and the index of refraction in GGX. In (b) we report the relation between the roughness in the Ward model and the roughness in GGX, for 3 values of the specular reflectance in Ward (or equivalently, the index of refraction in GGX, given (a)). The asterisks refer to the output of the genetic algorithm, the solid line the fitted values.

We repeated the experiment using a set of different diffuse colours for the spheres, but with the same incident lighting in all cases (compare the first and the second rows of Figure 5.9). Some of the images generated by this experiment, along with the average NRMSD are reported in Figure 5.16; the per channel NRMSD is reported in table 5.4. Overall, the measured NRMSD is about 13.37% lower than in the previously described unconstrained experiment. We decided to include the images generated for this experiment in a user study, to check whether such improvement in the NRMSD also corresponds to a better evaluation by the participants than in the previous case.

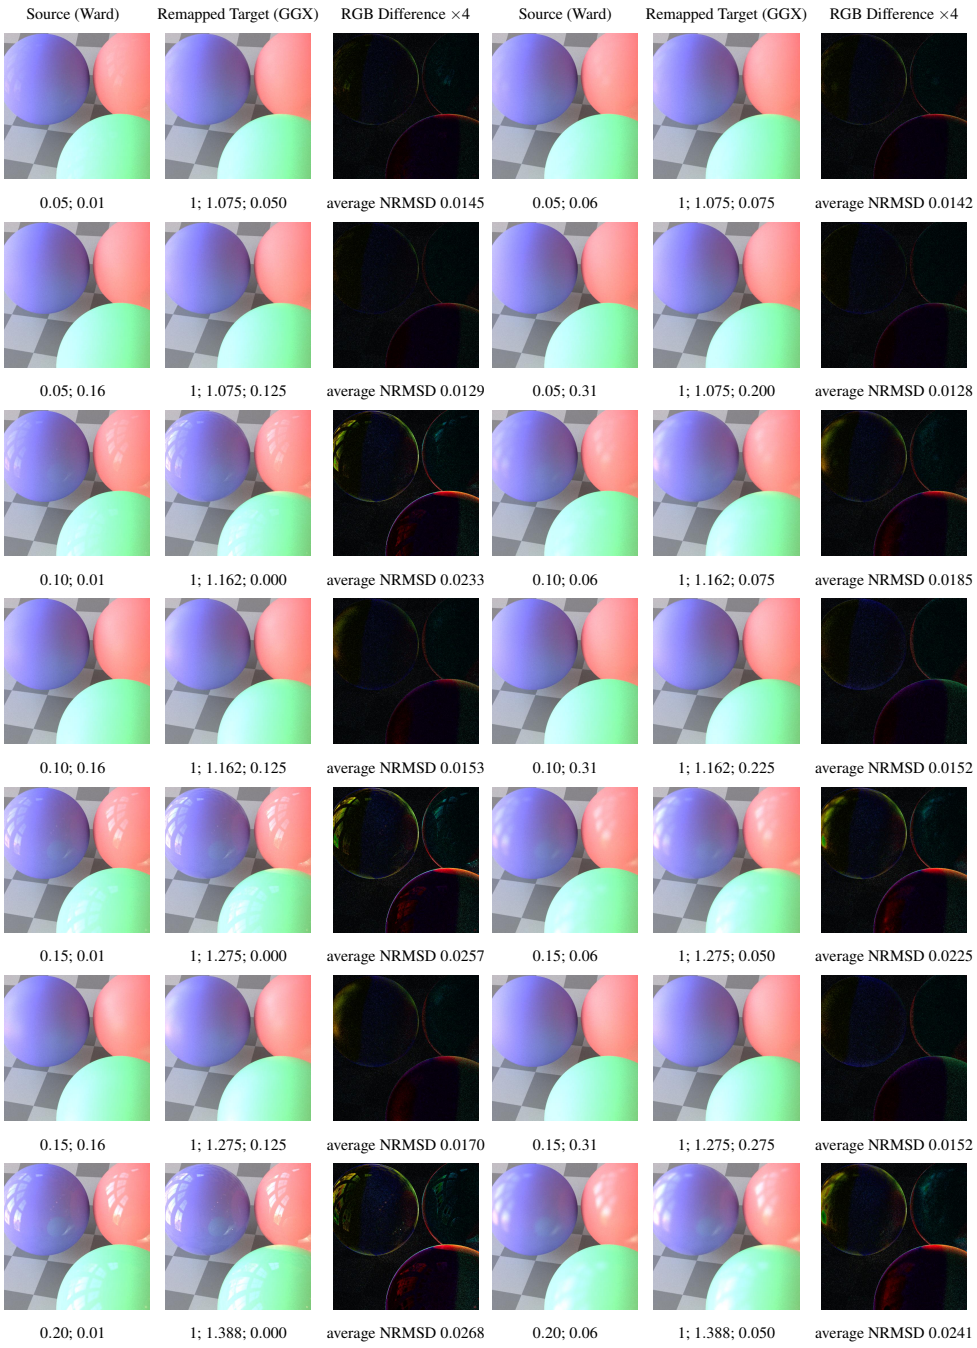


Figure 5.16: Remapping of dielectrics using prior knowledge of the functional relationships among parameters in the source and target models. We show the source images (first and forth columns), the corresponding remapped target images (second and fifth columns) and the absolute differences in the RGB colour space, amplified by a factor of 4 for visualisation purposes (third and sixth columns). Under each of the source images we report the values of the parameters $spec_{ward}$ and $alphaU_{ward} = alphaV_{ward}$ used to render it; under the target images we report the values of the parameters $spec_{ggx}$, $intIOR_{ggx}$ and $alpha_{ggx}$ found by our framework. Finally, under each difference image we report the average NRMSD over the RGB color channels.

Table 5.4: Remapping of dielectrics using prior knowledge of the functional relationships among parameters in the source and target models. The first column reports the $spec_{ward}$ and $alphaU_{ward} = alphaV_{ward}$ parameters of the source images, the secondo column reports the corresponding $spec_{ggx}$, $intIOR_{ggx}$ and $alpha_{ggx}$ parameters found by our framework. The last three columns report the Normalised Root Mean Square Deviation, per color channel.

| Source BRDF | Target BRDF | NRMSD Red | NRMSD Green | NRMSD Blue |
|---|---|---|---|---|
| (0.00 , 0.01) | (0 , 1.013 , 0.025) | 0.0102 | 0.0119 | 0.0160 |
| (0.05 , 0.01) | (1 , 1.075 , 0.050) | 0.0138 | 0.0134 | 0.0164 |
| (0.05 , 0.06) | (1 , 1.075 , 0.075) | 0.0135 | 0.0130 | 0.0160 |
| (0.05 , 0.16) | (1 , 1.075 , 0.125) | 0.0111 | 0.0118 | 0.0156 |
| (0.05 , 0.31) | (1 , 1.075 , 0.200) | 0.0104 | 0.0118 | 0.0161 |
| (0.05 , 0.71) | (1 , 1.075 , 0.400) | 0.0114 | 0.0128 | 0.0176 |
| (0.10 , 0.01) | (1 , 1.162 , 0.000) | 0.0260 | 0.0228 | 0.0211 |
| (0.10 , 0.06) | (1 , 1.162 , 0.075) | 0.0203 | 0.0170 | 0.0181 |
| (0.10 , 0.16) | (1 , 1.162 , 0.125) | 0.0148 | 0.0141 | 0.0169 |
| (0.10 , 0.31) | (1 , 1.162 , 0.225) | 0.0122 | 0.0139 | 0.0194 |
| (0.10 , 0.71) | (1 , 1.162 , 0.438) | 0.0178 | 0.0192 | 0.0268 |
| (0.15 , 0.01) | (1 , 1.275 , 0.000) | 0.0289 | 0.0251 | 0.0232 |
| (0.15 , 0.06) | (1 , 1.275 , 0.050) | 0.0256 | 0.0217 | 0.0203 |
| (0.15 , 0.16) | (1 , 1.275 , 0.125) | 0.0176 | 0.0160 | 0.0175 |
| (0.15 , 0.31) | (1 , 1.275 , 0.275) | 0.0120 | 0.0140 | 0.0196 |
| (0.15 , 0.71) | (1 , 1.275 , 0.523) | 0.0219 | 0.0234 | 0.0327 |
| (0.20 , 0.01) | (1 , 1.388 , 0.000) | 0.0304 | 0.0255 | 0.0246 |
| (0.20 , 0.06) | (1 , 1.388 , 0.050) | 0.0276 | 0.0232 | 0.0216 |
| (0.20 , 0.16) | (1 , 1.388 , 0.150) | 0.0197 | 0.0179 | 0.0180 |
| (0.20 , 0.31) | (1 , 1.388 , 0.275) | 0.0123 | 0.0140 | 0.0187 |
| (0.20 , 0.71) | (1 , 1.388 , 0.527) | 0.0187 | 0.0211 | 0.0300 |
| (0.25 , 0.01) | (1 , 1.525 , 0.000) | 0.0303 | 0.0253 | 0.0253 |
| (0.25 , 0.06) | (1 , 1.525 , 0.050) | 0.0272 | 0.0226 | 0.0219 |
| (0.25 , 0.16) | (1 , 1.525 , 0.150) | 0.0198 | 0.0184 | 0.0181 |
| (0.25 , 0.31) | (1 , 1.525 , 0.275) | 0.0139 | 0.0145 | 0.0177 |
| (0.25 , 0.71) | (1 , 1.525 , 0.592) | 0.0170 | 0.0198 | 0.0283 |
| (0.30 , 0.01) | (1 , 1.663 , 0.000) | 0.0292 | 0.0257 | 0.0274 |
| (0.30 , 0.06) | (1 , 1.663 , 0.050) | 0.0266 | 0.0223 | 0.0221 |
| (0.30 , 0.16) | (1 , 1.663 , 0.125) | 0.0216 | 0.0196 | 0.0192 |
| (0.30 , 0.31) | (1 , 1.663 , 0.275) | 0.0143 | 0.0149 | 0.0172 |
| (0.30 , 0.71) | (1 , 1.663 , 0.629) | 0.0158 | 0.0189 | 0.0267 |
| (0.35 , 0.01) | (1 , 1.838 , 0.000) | 0.0276 | 0.0237 | 0.0249 |
| (0.35 , 0.06) | (1 , 1.838 , 0.025) | 0.0282 | 0.0249 | 0.0251 |
| (0.35 , 0.16) | (1 , 1.838 , 0.125) | 0.0186 | 0.0172 | 0.0184 |
| (0.35 , 0.31) | (1 , 1.838 , 0.275) | 0.0137 | 0.0146 | 0.0168 |
| (0.35 , 0.71) | (1 , 1.838 , 0.629) | 0.0155 | 0.0183 | 0.0257 |

### 5.5.2 Numerical Validation - Conductors

The experiments on conductors remapping involved a remapping from the source Ashikhmin-Shirley model to the target Cook-Torrance, which implies a mapping $\{alphaU_{as} = alphaV_{as}\} \rightarrow \{alpha_{ct}\}$, hence the dimensionality of the remapping is lower than the dielectric case. In Figure 5.18 we report the output of the fitting module, which nicely summarises the relation between the roughness in the two models with a $2^{nd}$ degree polynomial; for values of roughness smaller than 0.2 the remapped values are coincident with the source values, whereas for higher values of the source roughness the remapped roughness differs noticeably. Table 5.5 reports the per channel NRMSD among source and target models, for each of the test images; 16 out of the 19 image pairs used for this experiment are reported in Figure 5.17.

We repeated our experiment using this time the GGX model as target; while our framework managed to provide a visually plausible remapping from the source to the target parameters (see Figure 5.19), the average NRMSD is up to 5.5 bigger than in the previous case, as reported in table 5.5. This is due to the very different characteristics of the microfacets distribution at the core of the GGX model, which is engineered to have a narrow specular peak and a tail much longer than usual. As a consequence, a rendering using a GGX model displays much more shadowing than most current models [5] and such effect simply cannot be reproduced. Figure 5.21, first column, reports three renderings of an iridium sphere using the Ashikhmin-Shirley model, for 3 different values of surface roughness; the second and third columns reports the corresponding renderings with the remapped parameters, respectively with the Cook-Torrance and the GGX model. The effect of the aforementioned tail is to spread out the incoming light over a much wider area than a Gaussian distribution, around a narrower specular peak.

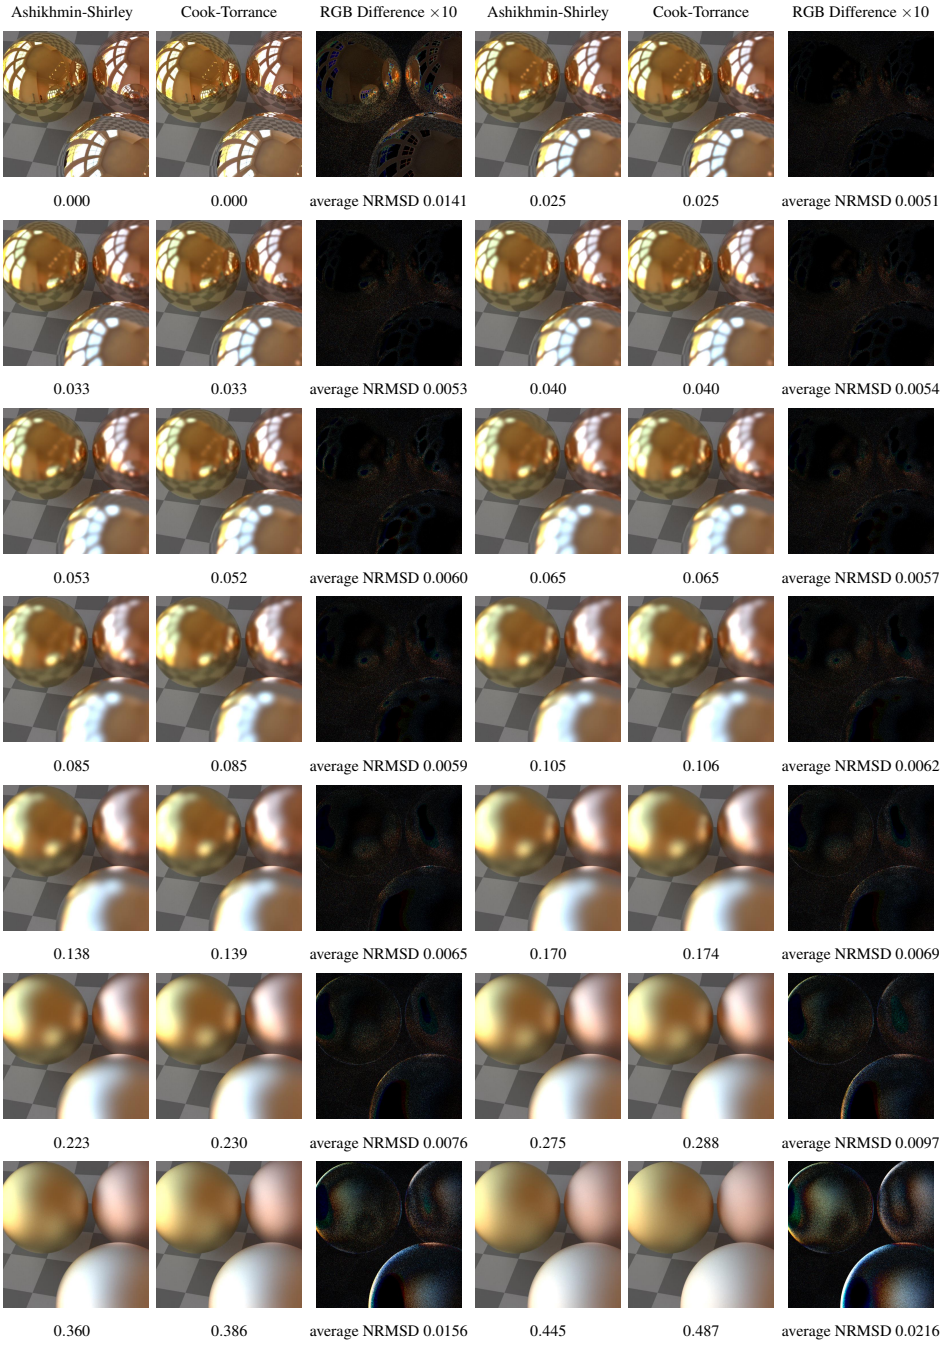


Figure 5.17: Remapping of conductors, remapping from Ashikhmin-Shirley to Cook-Torrance. We show the source images (first and forth columns), the corresponding remapped target images (second and fifth columns) and the absolute differences in the RGB colour space, amplified by a factor of 10 for visualisation purposes (third and sixth columns). Under each of the source images we report the values of the parameters $alphaU_{as} = alphaV_{as}$ used to render it; under the target images we report the values of the parameters $alpha_{ct}$ found by our framework in the unconstrained setting. Finally, under each difference image we report the average NRMSD over the RGB color channels. The metals rendered in the images are gold, silver and copper, whereas the remapping was learned using iridium.

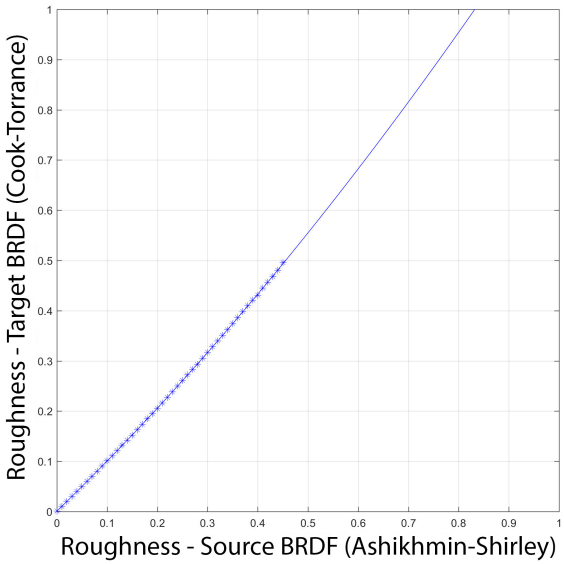


Figure 5.18: Output of the fitting module, experiment with conductors (Ashikhmin-Shirley → Cook-Torrance). Relation between the roughness in the source and target models. The asterisks refer to the output of the genetic algorithm, the solid line the fitted values.

Table 5.5: Remapping of conductors, Ashikhmin-Shirley → Cook-Torrance. The first column reports the $alphaU_{as} = alphaV_{as}$ parameter of the source images, the second column reports the corresponding $alpha_{ct}$ parameter found by our framework. The last three columns report the Normalised Root Mean Square Deviation, per color channel.

| Source BRDF | Target BRDF | NRMSD Red | NRMSD Green | NRMSD Blue |
|---|---|---|---|---|
| 0.000 | 0.000 | 0.0180 | 0.0137 | 0.0105 |
| 0.008 | 0.008 | 0.0054 | 0.0046 | 0.0043 |
| 0.015 | 0.015 | 0.0055 | 0.0048 | 0.0045 |
| 0.020 | 0.020 | 0.0056 | 0.0049 | 0.0046 |
| 0.025 | 0.025 | 0.0057 | 0.0050 | 0.0047 |
| 0.033 | 0.033 | 0.0059 | 0.0052 | 0.0049 |
| 0.040 | 0.040 | 0.0060 | 0.0053 | 0.0050 |
| 0.053 | 0.052 | 0.0064 | 0.0059 | 0.0057 |
| 0.065 | 0.065 | 0.0062 | 0.0056 | 0.0053 |
| 0.085 | 0.085 | 0.0064 | 0.0058 | 0.0056 |
| 0.105 | 0.106 | 0.0066 | 0.0061 | 0.0059 |
| 0.138 | 0.139 | 0.0069 | 0.0063 | 0.0062 |
| 0.170 | 0.174 | 0.0072 | 0.0068 | 0.0068 |
| 0.223 | 0.230 | 0.0079 | 0.0075 | 0.0075 |
| 0.275 | 0.288 | 0.0100 | 0.0097 | 0.0094 |
| 0.360 | 0.386 | 0.0158 | 0.0156 | 0.0155 |
| 0.445 | 0.487 | 0.0211 | 0.0216 | 0.0220 |
| 0.583 | 0.660 | 0.0269 | 0.0286 | 0.0294 |
| 0.720 | 0.843 | 0.0305 | 0.0320 | 0.0329 |

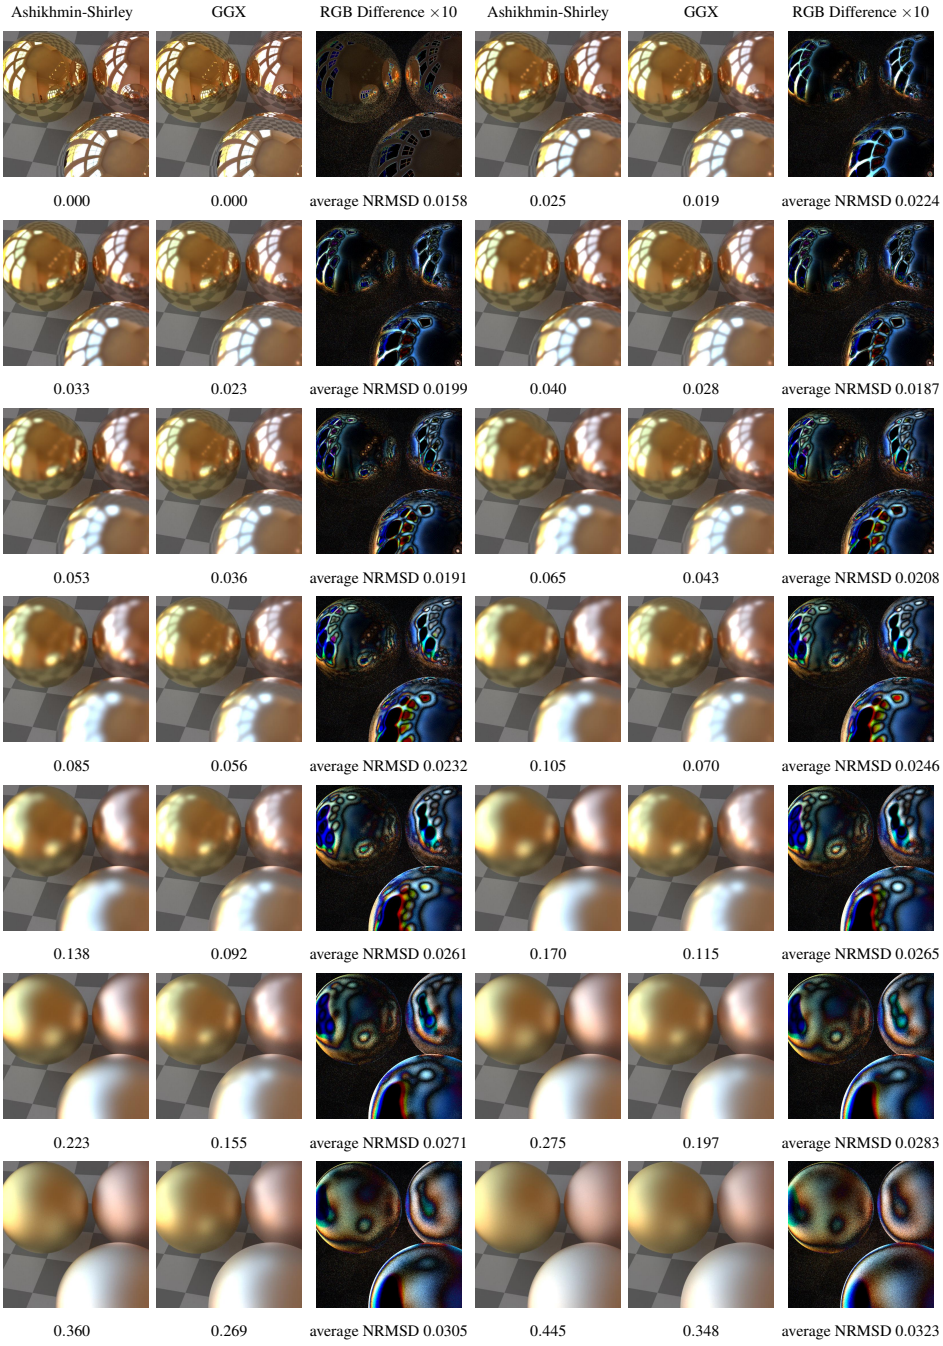


Figure 5.19: Remapping of conductors, remapping from Ashikhmin-Shirley to GGX. We show the source images (first and forth columns), the corresponding remapped target images (second and fifth columns) and the absolute differences in the RGB colour space, amplified by a factor of 10 for visualisation purposes (third and sixth columns). Under each of the source images we report the values of the parameters $alphaU_{as} = alphaV_{as}$ used to render it; under the target images we report the values of the parameters $alpha_{ggx}$ found by our framework in the unconstrained setting. Finally, under each difference image we report the average NRMSD over the RGB color channels. The metals rendered in the images are gold, silver and copper, whereas the remapping was learned using iridium.

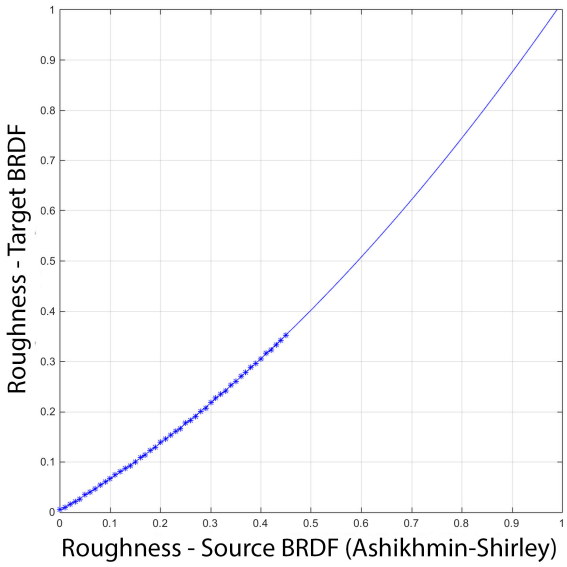


Figure 5.20: Output of the fitting module, experiment with conductors (Ashikhmin-Shirley → GGX). Relation between the roughness in the source model and the roughness in the target model. The asterisks refer to the output of the genetic algorithm, the solid line the fitted values.

Table 5.6: Remapping of conductors, Ashikhmin-Shirley → GGX. The first column reports the $alphaU_{as} = alphaV_{as}$ parameter of the source images, the second column reports the corresponding $alpha_{ct}$ parameter found by our framework. The last three columns report the Normalised Root Mean Square Deviation, per color channel.

| Source BRDF | Target BRDF | NRMSD Red | NRMSD Green | NRMSD Blue |
|---|---|---|---|---|
| 0.000 | 0.000 | 0.0197 | 0.0154 | 0.0123 |
| 0.008 | 0.008 | 0.0214 | 0.0266 | 0.0315 |
| 0.015 | 0.012 | 0.0199 | 0.0248 | 0.0298 |
| 0.020 | 0.016 | 0.0194 | 0.0237 | 0.0279 |
| 0.025 | 0.019 | 0.0189 | 0.0224 | 0.0258 |
| 0.033 | 0.023 | 0.0176 | 0.0198 | 0.0225 |
| 0.040 | 0.028 | 0.0166 | 0.0183 | 0.0212 |
| 0.053 | 0.036 | 0.0168 | 0.0187 | 0.0219 |
| 0.065 | 0.043 | 0.0180 | 0.0205 | 0.0239 |
| 0.085 | 0.056 | 0.0199 | 0.0229 | 0.0267 |
| 0.105 | 0.070 | 0.0216 | 0.0243 | 0.0279 |
| 0.138 | 0.092 | 0.0231 | 0.0256 | 0.0295 |
| 0.170 | 0.115 | 0.0237 | 0.0257 | 0.0302 |
| 0.223 | 0.155 | 0.0248 | 0.0261 | 0.0303 |
| 0.275 | 0.197 | 0.0267 | 0.0276 | 0.0306 |
| 0.360 | 0.269 | 0.0302 | 0.0301 | 0.0313 |
| 0.445 | 0.348 | 0.0335 | 0.0321 | 0.0314 |
| 0.583 | 0.490 | 0.0421 | 0.0387 | 0.0337 |
| 0.720 | 0.646 | 0.0546 | 0.0483 | 0.0404 |

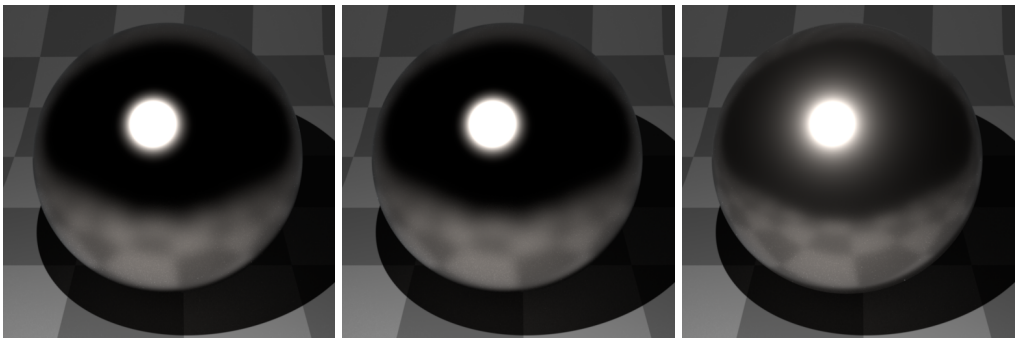

(a) Ashikhmin-Shirley, r: 0.105 (b) Cook-Torrance, r: 0.106 (c) GGX, r: 0.070

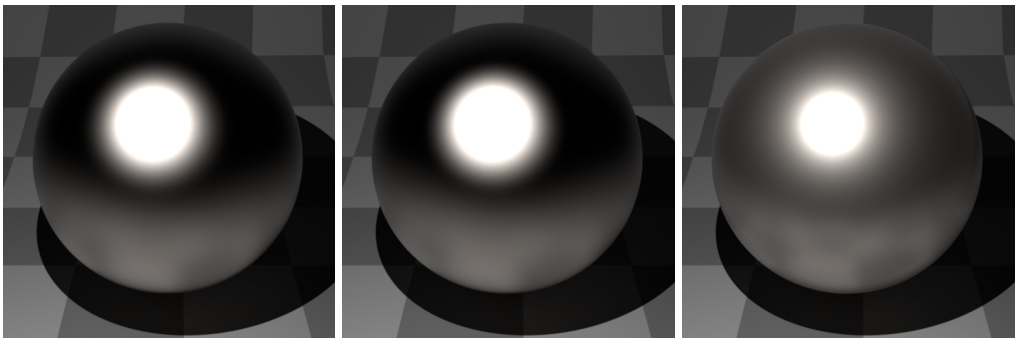

(d) Ashikhmin-Shirley, r: 0.223 (e) Cook-Torrance, r: 0.230 (f) GGX, r: 0.155

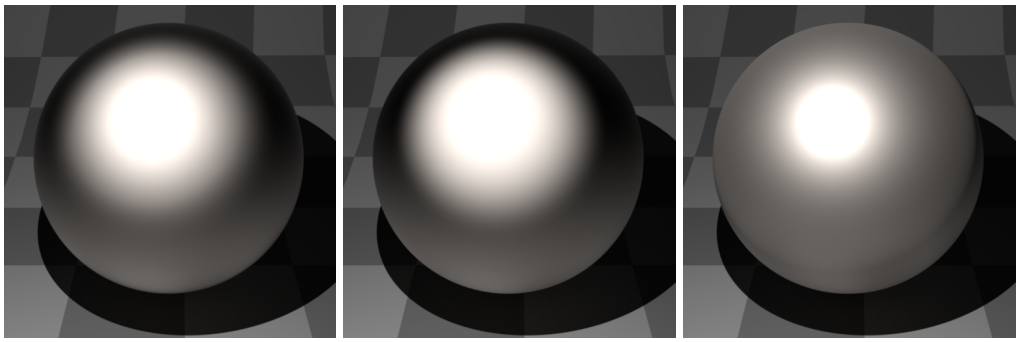

(g) Ashikhmin-Shirley, r: 0.445 (h) Cook-Torrance, r: 0.487 (i) GGX, r: 0.348

Figure 5.21: Visual comparison about the effect of the surface roughness among the Ashikhmin-Shirley, Cook-Torrance and GGX models. This is due to the very different characteristics of the microfacets distribution at the core of the GGX model, which is engineered to have a narrow specular peak and a tail much longer than usual. As a consequence, a rendering using a GGX model displays much more shadowing than most current models [5] and such effect simply cannot be reproduced.

## 5.6 User Study

The numerical validation provides an initial assessment of the quality of the remapping performed by our framework. However, in order to ensure that the remapping produces visually plausible results, regardless the selected metric, results were also verified by means of user studies, a common way to validate the results in the Computer Graphics field. Available options for user studies includes single-stimulus, double-stimulus, forced-choice pairwise comparison, and similarity judgements, which have been compared by Mantiuk *et al.* [160]. They found that one of the most accurate and efficient way to validate a metric is the forced-choice pairwise comparison, which we selected for our experiments, similarly to Glencross *et al.* [161] and Ward-Glencross [162].

The goal of the user validation studies is to reliably establish if the participants can distinguish between the appearance of an image rendered using the source BRDF and the same image rendered using the target BRDF, where the parameters are derived by our remapping framework. In our forced chose pairwise comparison design participant were presented with a sequence of two images next to each other (stimuli), shown on a calibrated display.

Each image pair consists of a complex scene rendered using the source BRDF model and the same scene rendered with the target BRDF using the remapped parameters. Each participant, after observing an image pair for exactly 3 seconds, were asked to rate the similarity between the displayed images, using integer ratings on a scale $0 - 2$, where 0 means "noticeably different", 1 encodes "slightly different" and 2 means "same"; the instructions were provided before starting the test. The order in which the image pairs is shown to the subject is randomised and includes a set of image pairs in which the images are actually the same, in order to provide a reference for a statistical analysis.

Every participant filled in the post study questionnaire stating their gender, age, vision problems, profession and level of expertise in graphics and art, along with their subjective perception about the difficulty in finding differences within the image pairs; this demographic information was used to profile the participants in order to identify any bias in the study [162].

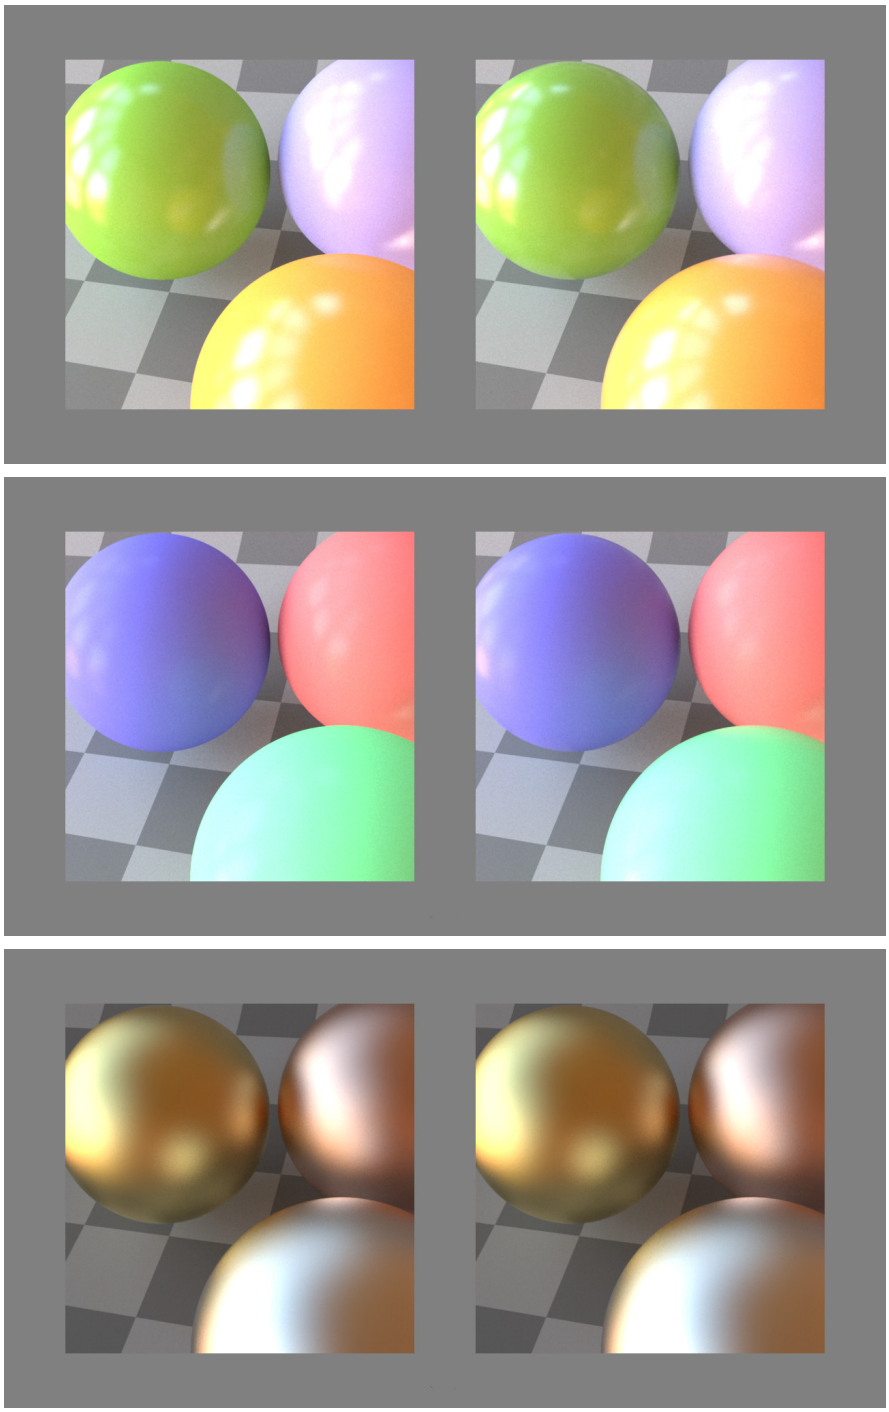

Figure 5.22: Some examples of source and remapped image pairs used in our user study. The first row shows the scene rendered with the source BRDF, the second column shows the corresponding remapped target BRDF. The first row refers to the experiment with dielectrics, in the unconstrained setting. The second rows refers to the second experiment with dielectrics, using prior knowledge of the semantic relationships among parameters. The bottom row refers to the experiment with metals.

Two different kinds of test scene have been used. The first kind of test scene consists of three spheres close to each other, with a different colour each (due to a different diffuse reflectance in case of dielectrics remapping, different metal in case of conductors remapping) but same BRDF parameters; the spheres are placed over a checkerboard and illuminated by environment lighting with a rich visual content and different from the one used by the genetic algorithm to remap the parameters (see Figure 5.22). The second kind of test scene makes use of three complex shapes out of the ones available from the Stanford 3D Scanning Repository, placed inside a Cornell box.

Overall four user studies were conducted:

Study 1) We presented dielectric materials with varying roughness and specularity, in order to validate our unconstrained remapping technique described in Section 5.4.1. The source model was the Ward BRDF model, whereas the target was the GGX model (an example of image pair is reported in Figure 5.22, top row). The study consisted of 36 test pairs and 5 control pairs.

Study 2) The second user study presents dielectric materials with varying roughness and specularity, aimed to validate our framework in the configuration using prior knowledge, described in Section 5.4.1. The source and target BRDFs were the same than in the previous user study. The diffuse colours of the spheres were different from the previous test, in order to confirm some findings obtained from it, as described in the following. An example of the image pairs used for this study is reported in Figure 5.22, middle row. Similarly to the first user study, we presented the participants with 36 test pairs and 5 control image pairs.

Study 3) The third user study presented a set of conductors (metals) with varying roughness, where the remapping from the source (Ashikhmin-Shirley) to the target model (Cook-Torrance) was performed with the technique based on prior knowledge (section 5.4.2); in Figure 5.22, bottom row, one of the stimuli is reported. A total of 19 test pairs, plus 2 control pairs, were used.

Study 4) The fourth and last user study was based on the second kind of test scene, where the selected shapes are the “Stanford Bunny”, “Happy Buddha” and

"Dragon". For each test pair, the source image was generated by randomly assigning to the objects one of the source BRDF models Ward or Ashikhmin-Shirley, with parameters randomly selected as well; the target image was generated by assigning to each object the corresponding target BRDF model and remapped parameters. The participants were presented with 2 test pairs and 2 control pairs.

The participants age range is between 14 and 71 years old and the level of expertise in graphics and arts varies from no expertise to professional. Since the user study 2 and 4 were performed together, at a different time than user study 1 and 3 (the last two performed in the same session), the participants to the user study 2 and 4 (13 subjects) are different from the ones which took part in the other two studies (15 subjects), with the exception of participant 12. In total there were 27 different participants and the demographic composition of the 2 groups of 15 and 13 participants were very similar.

As mentioned before, the participants were asked to indicate the perceived difficulty in finding differences within an image pair. The results of the questionnaire shows that in most of the images the subjects had difficulties in spotting differences, being most of them hardly noticeable. As show in Figure 5.23, in many cases participants gave similar average scores to the original-remapped image pairs (test pairs) and to the same-image pairs (control pairs), in some cases even higher. Participant 12 performed particularly well in the task of distinguishing among test and control pairs, giving consistently higher scores to the latter than to the test ones. In the post study questionnaire participant 12 reported to be a professional designer and theatrical makeup artist, with no vision impairment.

The outcome of the user studies has been analyzed for statistical significance using the one-tailed $t-test$ and Welch's $t-test$ accounting for both the possibilities that the variances of the test and control pairs is equal or unequal. The null hypothesis is that the average rating of the test pairs is not significantly different from the average rating of the control pairs, and hence the observer cannot distinguish between the source and the remapped images; the alternative hypothesis is that the average rating of the control group is significantly greater than the average rating of the test pairs. All the tests were conducted at a significance level $\alpha = 0.05$.

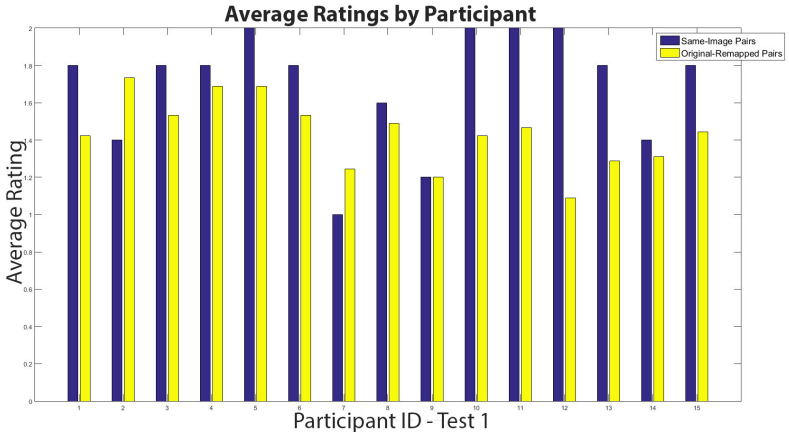


(a) User Study 1 - Average Ratings by Participant

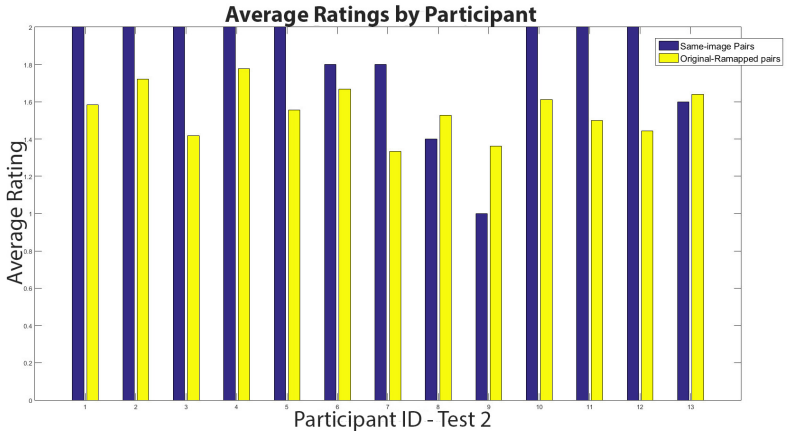


(b) User Study 2 - Average Ratings by Participant

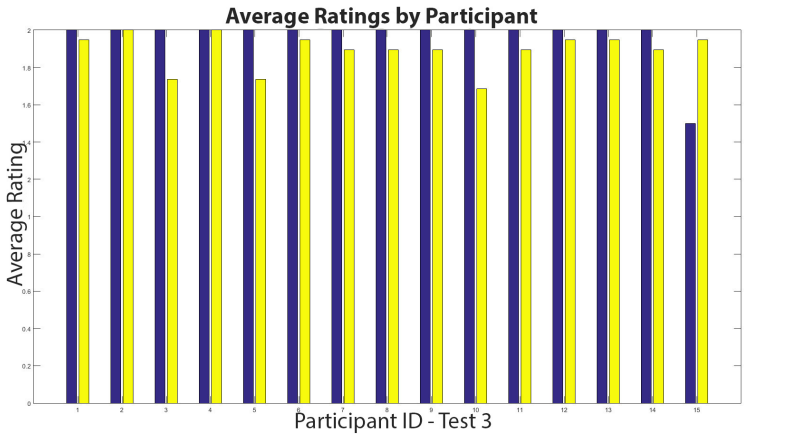


(c) User Study 3 - Average Ratings by Participant

Figure 5.23: Average Ratings by participant for each of the 3 user studies. The yellow bars refer to the ratings for the test pairs (original-remapped), whereas the blue bars refer to the control pairs (same-image pairs).

*User Study 1.* The average rating for the test pairs was 1.4296 and the average rating for the control pairs 1.6933; in Figure 5.24(a) the average ratings by image pair is reported. To test our results for statistical significance the $t-test$ (assuming equal variance between the two groups of image pairs) and Welch's $t-test$ (assuming unequal variance) were used (2 groups only). In both cases the t-statistic and p-value showed a significant main effect of pair-type and reject the null-hypothesis, being $t = 3.78$ (degrees of freedom $df = 613$), $p = 0.00$ for the first test and $t = 4.21$ (degrees of freedom $df = 103.41$), $p = 0.00$ for the second one.

In order to highlight specific areas of the source BRDF parameters space in which the effectiveness of the remapping could be lower, we assigned each test pair to a different group and repeated the test. The pairwise comparison results from a multiple comparison test using the Tukey's Honestly Significant Difference Procedure, conducting the hypothesis tests at the 5% significance level, allowed to highlight that the participants tend to spot differences between source and remapped image pairs when $0 < alpha_{ward} \leq 0.06$, or when $spec_{ward} \geq 0.35$. Hence, it is possible to partition the Ward's parameters space in two areas, as shown in figure 5.25. The red area corresponds to parameters for which the remapping leads to unavoidable differences, in particular towards the silhouette of the objects, due to the Fresnel effect, modelled by the GGX model.

Accordingly, we partitioned the test image pairs into two groups:

- $red = \{\{0 \leq alpha_{ward} \leq 0.06\} \cap \{spec_{ward} \geq 0.35\}\}$,
- $green = \{\{alpha_{ward} > 0.06\} \cap \{spec_{ward} < 0.35\}\}$.

We repeated the statistical analysis comparing the mean of each group to the mean of the control group. Both the $t-test$ and the Welch's $t-test$ do not reject the null hypothesis when the green group is compared to the control one, being $t = 1.59$ (degrees of freedom $df = 358$), $p = 0.06$ for the first test and $t = 1.59$ (degrees of freedom $df = 116.32$), $p = 0.06$ for the second one, confirming that for renderings belonging to this large portion of the parameters space the observers were not able to reliably distinguish between source and remapped renderings. As for the red group, both tests reject the null hypothesis, $t = 5.83$ (degrees of freedom

$df = 328$), $p = 0.00$ and $t = 6.45$ (degrees of freedom $df = 143.05$), $p = 0.00$ respectively. The post-hoc analysis is reported in figure 5.29(b).

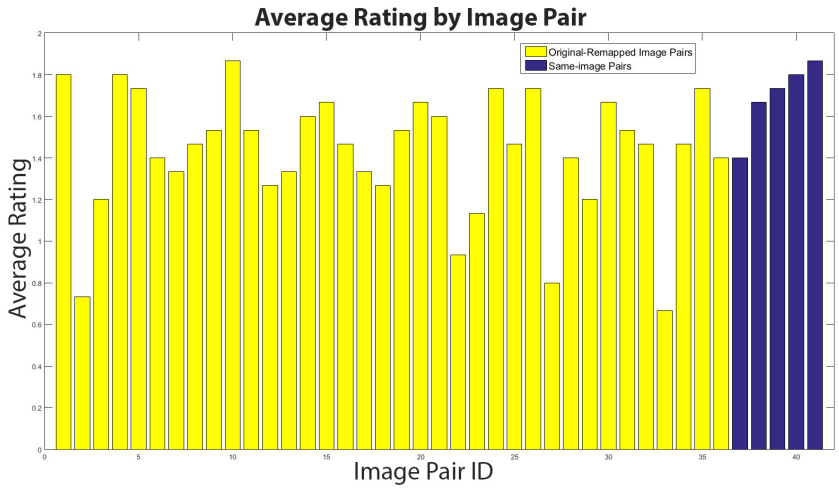


(a) Average Ratings by Image Pair

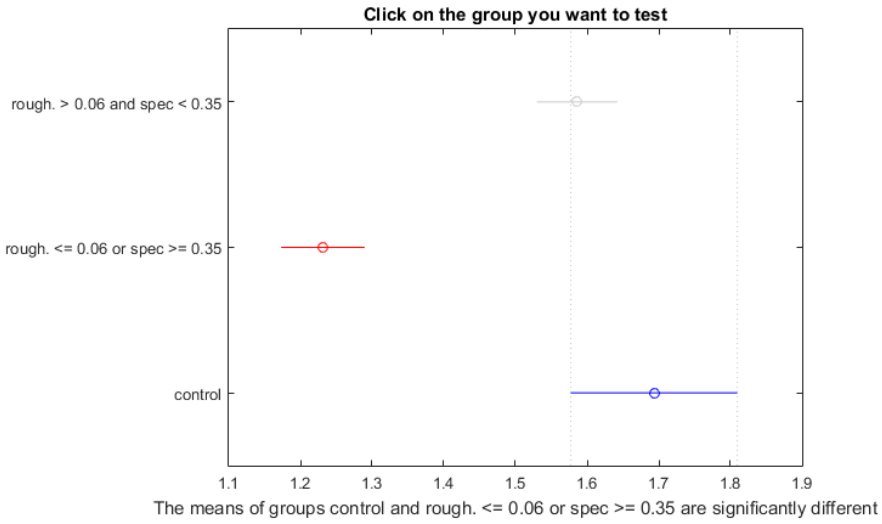


(b) Post hoc analysis

Figure 5.24: Statistical analysis of User Study 1.

We further investigated the results of the statistical analysis, by relating the average rating of each test pairs with the NRMSD between source and remapped image within each pair, as derived in section 5.5. In particular, we took into account the NRMSD on each of the RGB channels. To measure the statistical dependence we used the distance correlation [163], which allows to highlight a high

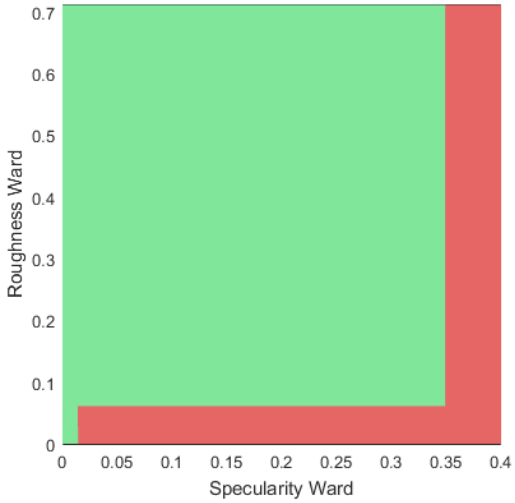


Figure 5.25: Partition of the Ward's parameters space according to the user studies. The red area refers to parameters for which the observers were able to distinguish between the source and target renderings.

correlation between the NRMSD on the blue channel and the observers ratings; figure 5.26 shows the average rating as a function of each of these measures.

A closer look at the images of the absolute differences in the RGB colour space between source and remapped images in each test pair reveals that most of the error is contained in the blue channel, and it is due to the reflection of the purple sphere over the edges of the green and orange spheres, caused by the Fresnel effect (Figure 5.27), not implemented by the Ward model. According to the participant perception, the Fresnel effect is more visible for low values of surface roughness and high values of the specular reflectance in the Ward model, the latter being remapped towards high values of the index of refraction in the GGX model.

*User Study 2.* The major source of error in the previous experiment was the reflection of the purple sphere (where the blue value dominates the red and green) over the other two spheres in the scene (figure 5.22, top row), which caused a higher NRMSD on the blue channel for low values of roughness and high values of specularity. In order to confirm the outcome of the first user study we decided to change the diffuse colours of the spheres in the scene (figure 5.22, middle row), in particular we replaced the purple sphere with a red one, hence this time we expected an higher correlation between the observers ratings and the NRMSD on

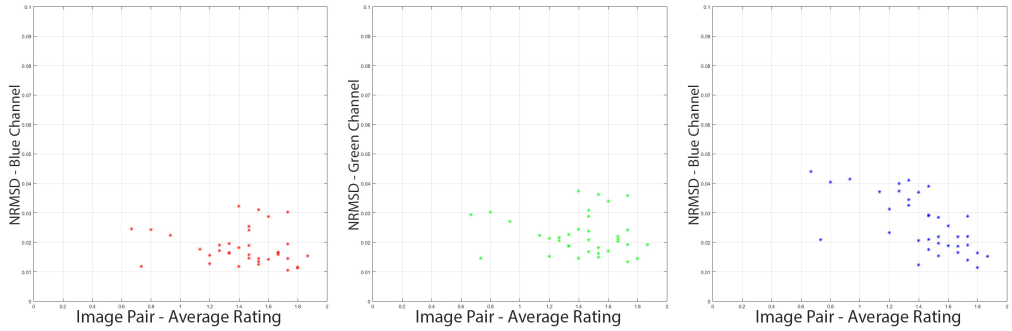


(a) Red Channel - Distance Correlation 0.3058 (b) Green Channel - Distance Correlation 0.2773 (c) Blue Channel - Distance Correlation 0.6531

Figure 5.26: Correlation between a numerical error measure (NRMSD) between source and remapped image within each test pair and the average rating by the participants to User Study 1. From left to right, the plots refer to NRMSD on the red, green and blue channels.

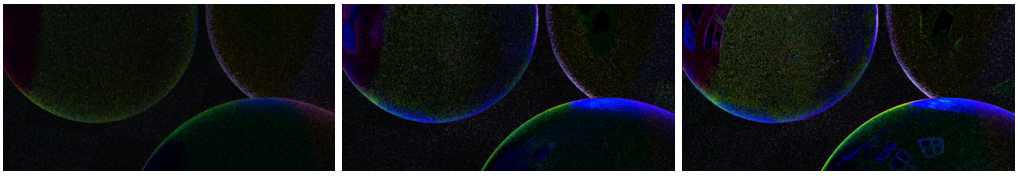

(a) NRMSD (blue): 0.015 ; average rating: 1.867 (b) NRMSD (blue): 0.037 ; average rating: 1.267 (c) NRMSD (blue): 0.042 ; average rating: 0.933

Figure 5.27: Correlation between the NRMSD over the blue channel and the average rating by the participants to User Study 1. From left to right the NRMSD increases, while the average rating decreases.

the red channel. Additionally, to generate the image pairs we used the constrained remapping strategy.

The average ratings per image pair are comparable to the ones in the first user study, with an average rating for the test pairs of 1.5491 and the average score for the control pairs 1.8154 (figure 5.29(a)). We carried out a set of statistical tests in the same way as done for the first user study, obtaining the following measures:

- All test pairs in one group. The $t-test$ (equal variance assumed for the control and test groups) rejects the null-hypothesis, with $t = 3.30$ (degrees of freedom $df = 531$), $p = 0.00$. Also the Welch's $t-test$ (unequal variances assumed) shows a significant main effect of pair type, rejecting the null hypothesis with $t = 4.6979$ (degrees of freedom $df = 117.18$), $p = 0.00$. Similarly to User study 1, the post hoc analysis allows to partition the Ward's parameters space into the same two regions reported in figure 5.25.

- Test image pairs partitioned into the groups *red* and *green*. Similarly to user study 1, both the $t-test$ and the Welch's $t-test$ reject the null-hypothesis when the image pairs belonging to the red group are compared to the control pairs, respectively with $t = 5.27$ (degrees of freedom $df = 284$), $p = 0.00$ and $t = 7.22$ (degrees of freedom $df = 201.57$), $p = 0.00$. As for the green group, both statistical tests confirm that for all source renderings in this portion of the parameters space the participants were not able to reliably distinguish between a source rendering and the corresponding remapped one, with $t = 0.97$ (degrees of freedom $df = 310$), $p = 0.17$ ($t-test$) and $t = 1.04$ (degrees of freedom $df = 110.36$), $p = 0.15$ (Welch's $t-test$). The post-hoc analysis is reported in figure 5.29(b).

The distance correlation was used again to investigate any statistical dependence of the average rating on the NRMSD. As expected, the test reported a high correlation with the NRMSD on the red channel, with a value of 0.6247; the values for the green and blue channels were respectively 0.5135 and 0.2984. In figure 5.28 we report a close up of the areas with the higher NRMSD in the red channel (compare with Figure 5.27).

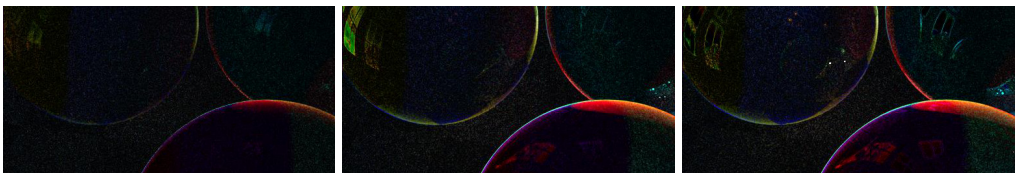

(a) NRMSD (red): 0.011 ; average rating: 2.000 (b) NRMSD (red): 0.014 ; average rating: 1.769 (c) NRMSD (red): 0.026 ; average rating: 1.077

Figure 5.28: Correlation between the NRMSD over the red channel and the average rating by the participants to User Study 2. From left to right the NRMSD increases, while the average rating decreases.

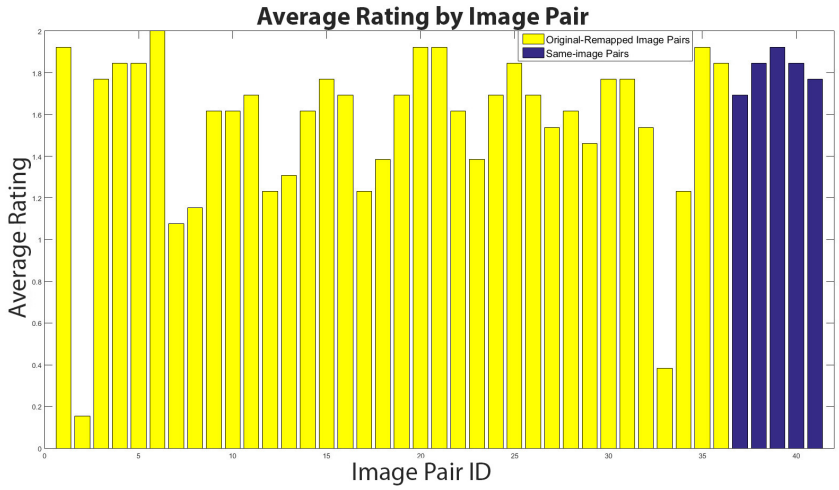


(a) Average Ratings by Image Pair

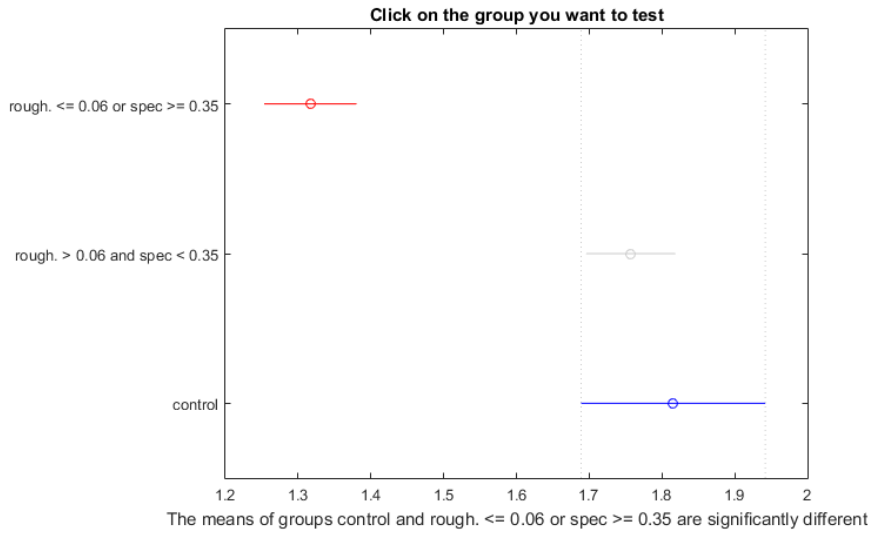


(b) Post hoc analysis red/green groups

Figure 5.29: Statistical analysis of User Study 2.

**User Study 3.** The third user study is meant to validate our remapping of conductors performed by our framework. Since the mapping was performed only in one dimension ($alpha_{as} \rightarrow alpha_{ct}$) the number of test and control pairs is lower than in the previous cases (respectively 19 and 2). The average rating for the test pairs was 1.8912, with an average rating for the control pairs of 1.9666; the average rating for all the images is reported in Figure 5.30. Both the $t-test$ and the Welch's $t-test$, performed by including all the test pairs in one group,

indicated no significant difference among test and control pairs, with $t = 1.08$ (degrees of freedom $df = 313$), $p = 0.14$ and $t = 1.54$ (degrees of freedom $df = 45.92$), $p = 0.07$. The outcome of the tests confirms the participant comments, which reported the impression that almost all the images were the same.

Also in this case we found a high correlation among user ratings and the NRMSD values, using the distance correlation measure: 0.7871, 0.7542 and 0.5579, respectively for the red, green and blue channels; the first two values are very similar, whereas the blue channel shows a lower correlation and this is likely due to the inclusion in the scene of a gold and a copper metal spheres, both having a low blue component.

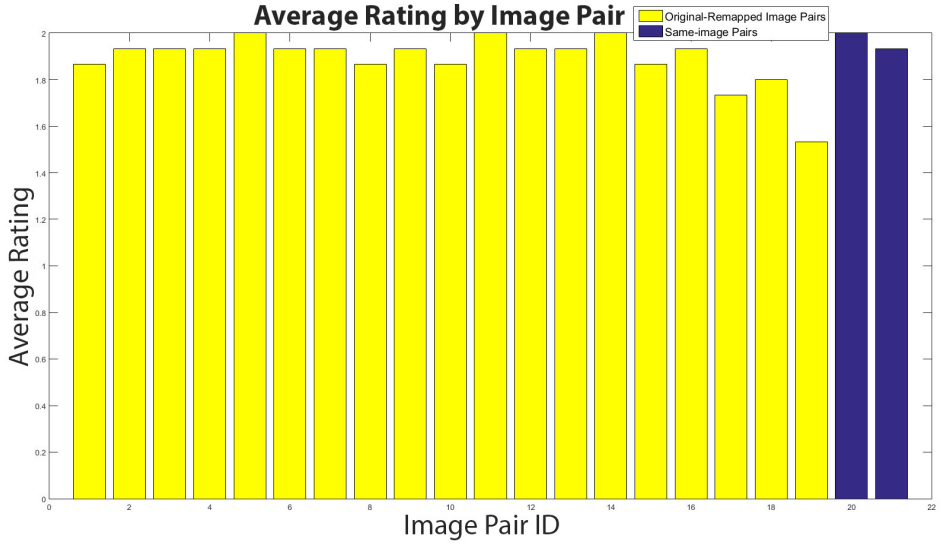


(a) Average Ratings by Image Pair

Figure 5.30: Statistical analysis of User Study 3.

**User Study 4.** The purpose of the fourth and last user study was to assess the effectiveness of our technique when combining several BRDF models and complex objects in one scene. The $t-test$, performed by including all the test pair in one group, indicated no significant difference among the average rating of test and control pairs, with $t = 1.54$ (degrees of freedom $df = 50$), $p = 0.07$ ($t-test$) and $t = 1.54$ (degrees of freedom $df = 42.24$), $p = 0.07$ (Welch's $t-test$). The average score for the class "test pair" was 1.7692 and 1.9231 for the class "control pair". The images used in this user study are reported in Figure 5.31.

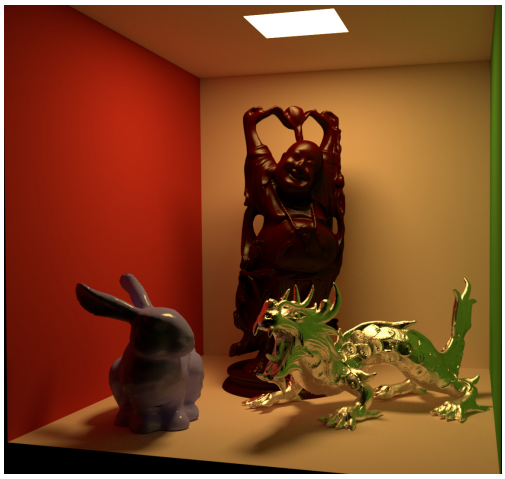

(a) Test pair 1 - source rendering

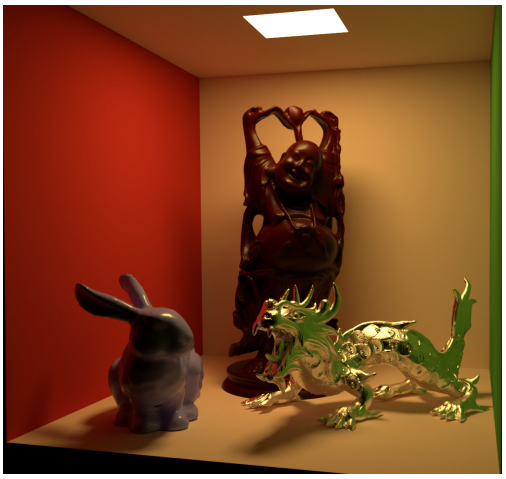

(b) Test pair 1 - target rendering

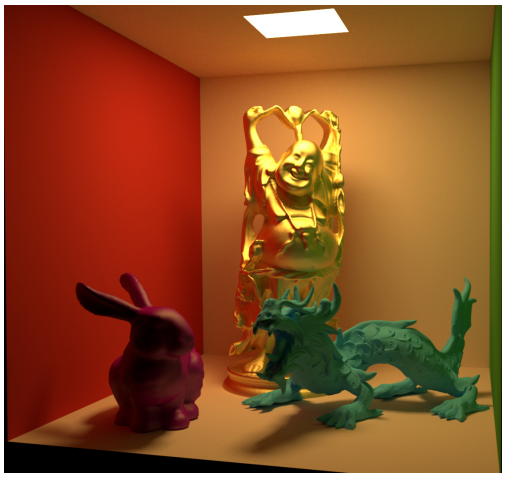

(c) Test pair 2 - source rendering

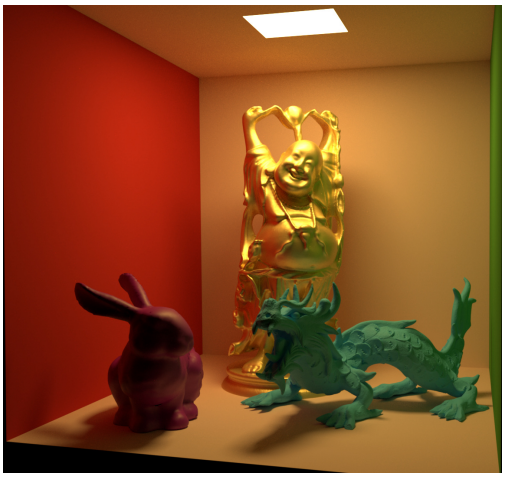

(d) Test pair 2 - target rendering

(e) Control pair 1 - source rendering

(f) Control pair 2 - source rendering

Figure 5.31: The set of images used for the test and control sets in User Study 4.

## 5.7 Psychometric scaling experiments

The user studies allowed to assess that the proposed framework provides, with both the constrained and unconstrained strategies, an effective tool to remap parameters from a given source model to the parameters of an arbitrary target one. Moreover, they highlighted that when there are noticeable differences between source and remapped renderings, this is due to deep differences between the characteristics of source and target models, such as in case of a remapping from the Ward to GGX (user studies 1 and 2). We further focus on this particular case and extend our validation with psychometric scaling experiments. The purpose of the additional experiments, with tools of the human perception field, is twofold:

- to provide an assessment that our framework provides not just a good remapping, but also it is the best possible, on the average, in its parameters neighborhood (*Perceptual Experiment 1*);
- to provide stronger evidence that what cause the remapped renderings in the GGX model to be distinguishable from source renderings with parameters in the red portion of the plane (figure 5.25), is actually the visually disruptive effect of the lack of the Fresnel effect on one side (Ward), and its presence on the other side (GGX) (*Perceptual Experiment 2*).

We designed our experiments in Matlab, using the Psychophysics Toolbox Version 3 extensions [164–166].

*Perceptual Experiment 1*. In the first experiment subjects are presented with the source rendering on the right side of the screen and nine renderings on the left side (figure 5.32). Only one rendering out of the 9 is given by our remapped parameters, whereas the others are rendered using parameters in the reflectance neighborhood of the target parameter space. We focused on the plane given by $intIOR_{ggx} \times alpha_{ggx}$, moving along one direction at a time.

The shapes are blobby object and their exact shape and order are randomised. The diffuse value is fixed to a green shade, with values $RGB = [0.1, 0.3, 0.1]$. The task of the observers is to select, on the left side, the closest rendering in terms of reflectance properties to the source rendering; once selected the image will disappear and the observer will then choose the closest one out the eight left and

so on, until all nine images have been removed from the screen. This process is repeated thirty times for each observer, for a total of 25 points in the Ward parameters space, spanning the whole range as specified in Section 5.4. The light probe used in this experiment is the "Courtyard of the Doge's palace, Venice, Italy" [167]. Overall 7 participant took part in this study, and each session lasted about 2 hours.

The rationale of the experiment is that on the average among repetitions and observers, the visually closer target rendering to the source one, corresponding to the best set of target parameters for that given source, will have a lower ranking; the ranking should increase with the distance of the neighbours to the best one, giving raise to a "V" shape centered on the latter, in a plot $neigbourID \times rank$.

The outcome of this experiment is reported in figures 5.33 and 5.34. For each of the plot on the $x-axis$ we report the neighbours identifiers, where 0 indicates the remapped parameters found by our framework (constrained strategy), of coordinates $\{I_0, R_0\}$ in the $intIOR_{ggx} \times alpha_{ggx}$ space; the coordinate of the neighbours $\{-4,-3,\ldots,0,\ldots,4\}$ are reported in table 5.7. Please note that some of the neighbours might not exist, which is the case when $\{I_0, R_0\}$ is close to the edges of the $intIOR_{ggx} \times alpha_{ggx}$ space. For instance, let $\{I_0, R_0\} = 1.21, 0.11$: in this case "$-3$" and "$-4$" in the $intIOR_{ggx}$ direction would have a value smaller than 1, hence we assign to them the same coordinates as "$-2$" (*i.e.* $\{1.01, 0.11\}$). Similar considerations apply to the $alpha_{ggx}$ direction.

The analysis of the neighbours in the $intIOR_{ggx}$ direction shows that in 56% of the cases the observer preferred the rendering(s) corresponding to the parameters remapped by our framework ("0"), and in 37.5% of the cases its immediate neighbour "$-1$", indicating a preference towards a slightly lower value of the Index of Refraction and hence a less visible Fresnel effect. As for the roughness direction, our remapped parameters were preferred again in the 56% of the cases and "$-1$" in 28% of the cases. Flat areas indicate regions of the target parameters space in which the renderings do not show perceivable differences with their immediate neighbours. In many cases the "$-1$'" neighbours have a very similar value of the fitness function (see figure 5.35), and the selection of $\{I_0, R_0\}$ by our framework is due to the constrained, monotonically non-decreasing strategy.

*Perceptual Experiment 2.* In the second experiment the observers are pre-

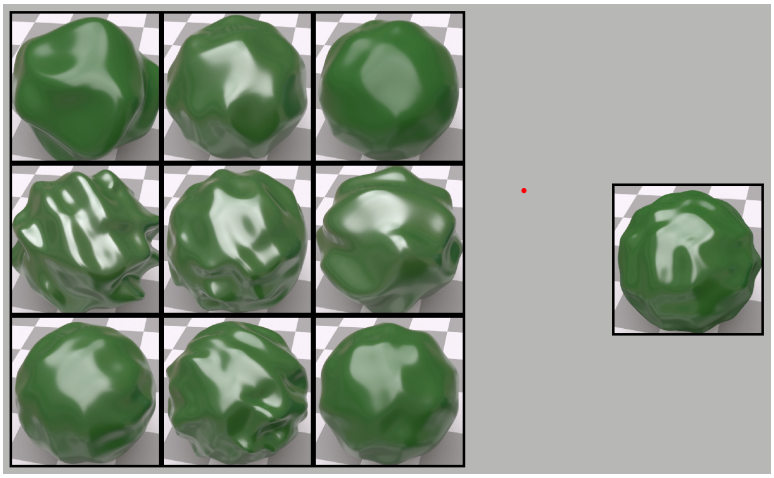

Figure 5.32: Interface of the first psychometric scaling experiment.

Table 5.7: Coordinates of the neighbours in the first psychometric experiment. $\{I_0, R_0\}$ indicates the coordinate of the remapped parameters found by our framework on the $intIOR_{ggx} \times alpha_{ggx}$ space.

| Neighbour ID | $intIOR_{ggx}$ direction | $aplha_{ggx}$ direction |
|---|---|---|
| -4 | $\{I_0 - 0.4, R_0\}$ | $\{I_0, R_0 - 0.2\}$ |
| -3 | $\{I_0 - 0.3, R_0\}$ | $\{I_0, R_0 - 0.15\}$ |
| -2 | $\{I_0 - 0.2, R_0\}$ | $\{I_0, R_0 - 0.1\}$ |
| -1 | $\{I_0 - 0.1, R_0\}$ | $\{I_0, R_0 - 0.05\}$ |
| 0 | $\{I_0, R_0\}$ | $\{I_0, R_0\}$ |
| 1 | $\{I_0 + 0.1, R_0\}$ | $\{I_0, R_0 + 0.05\}$ |
| 2 | $\{I_0 + 0.2, R_0\}$ | $\{I_0, R_0 + 0.1\}$ |
| 3 | $\{I_0 + 0.3, R_0\}$ | $\{I_0, R_0 + 0.15\}$ |
| 4 | $\{I_0 + 0.4, R_0\}$ | $\{I_0, R_0 + 0.2\}$ |

sented with a simpler interface, reported in figure 5.36. At the bottom of the interface a source BRDF rendering of a blobby shape is displayed, used as a reference, whereas at the top tow different renderings of blobby shapes are presented: one of them is rendered with the same model and parameters of the reference, while the second one is rendered using the remapped target parameters. In this forced-choice design the task of the user is to select, out of the two shapes at the top, the most similar to the reference in terms of reflectance properties. Similarly to *Perceptual Experiment 1*, we repeat this process thirty times for each observer and set of source parameters, for a total of 25 points in the Ward parameter space, spanning its whole range. The diffuse reflectance is set to $RGB = [0.1, 0.3, 0.1]$, while shape and left/right position is randomized for each repetition.

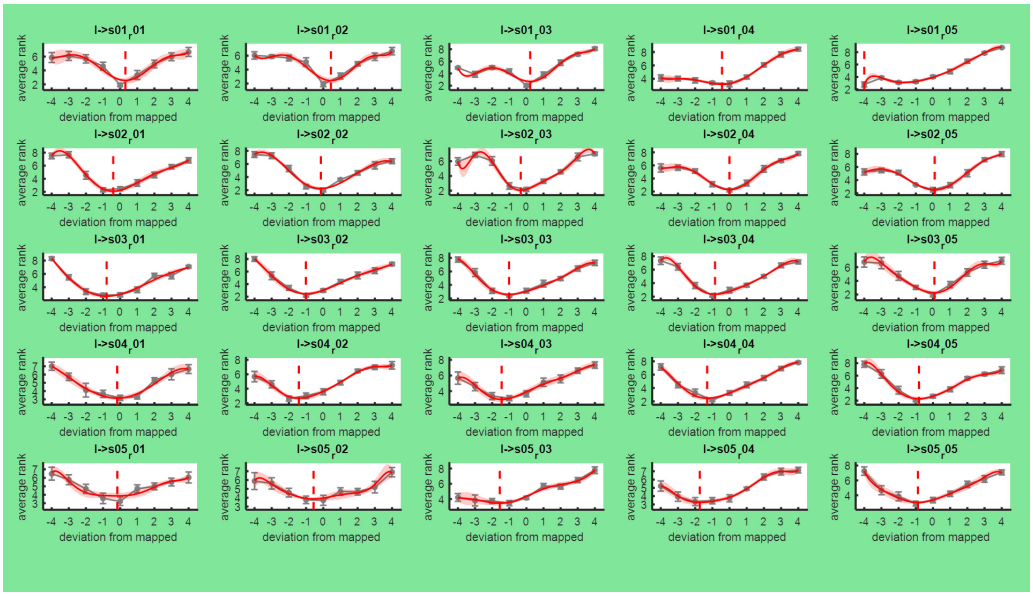


Figure 5.33: Results of the first experiment, neighbors in the Index of Refraction direction.

For each point in the Ward's parameter space, this experiments provides us with a percentage of times in which the observers correctly identified a source-source pair of blobby objects. In particular, a value of 100% indicates that the observers were always able to identify a target BRDF rendering from a source BRDF one, and a value of 50% indicates that the observers went randomly in the selection task, not being able to distinguish between a source and a remapped rendering.

We partitioned each blobby shape involved in the experiment in two areas, identified as "Fresnel" and "no Fresnel" by means of three different thresholds, respectively given by 102.5%, 110% and 120% of the reflectance at normal incidence (see figure 5.37). For each of this area, and for each corresponding shape, we calculated the average $\Delta E_{94}$ color distance [168] between corresponding source and target renderings, and related it to the percentage of correct answers by mean of logistic regression. As figure 5.38 shows, the slope of the curve related to the "no Fresnel" areas is almost flat, whereas the regression slope in the "Fresnel" areas clearly show a significant effect on the observers performance in distinguishing between source and remapped renderings, confirming the empirical finding from user studies 1 and 2: the observers actually use the Fresnel effect as a cue, mainly located towards the silhouette of the objects.

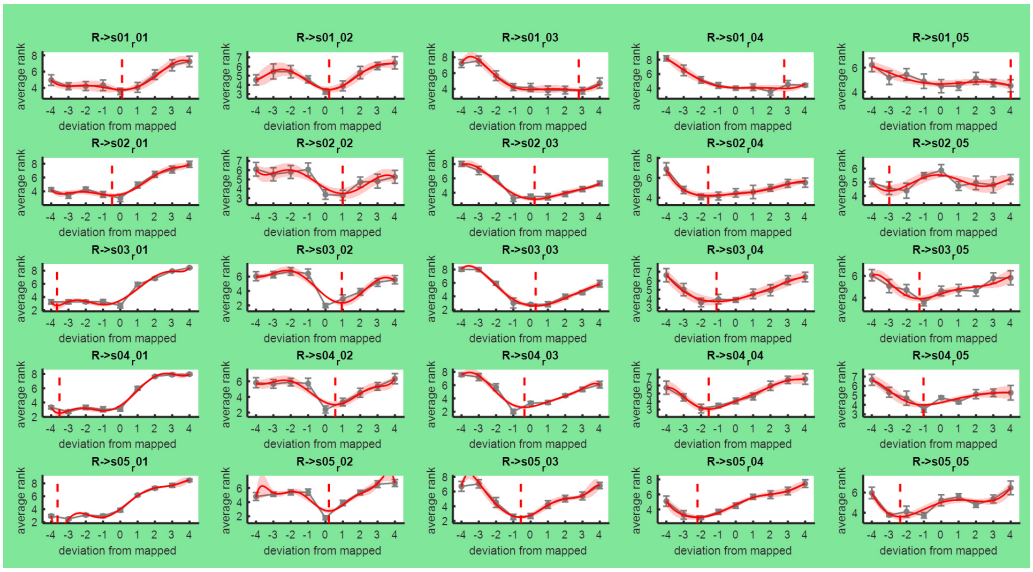


Figure 5.34: Results of the first experiment, neighbors in the roughness direction.

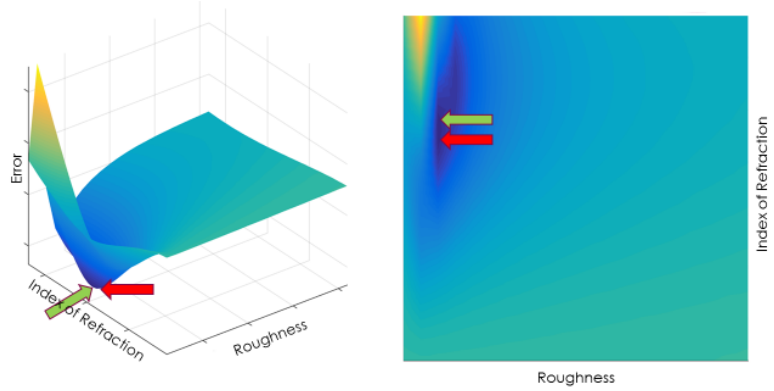


Figure 5.35: Error surface.

## 5.8 Conclusion

Very often a digital artist faces the problem of trying to match an image rendered somewhere else, with a BRDF model not available in the rendering tool in use. Standard practice is to manually tweak the parameters offered by the available model, based on a subjective evaluation and experience, in a time-consuming fashion. In this study we considered a “minimum-ignorance” scenario, in which the source material model (BRDF) and parameters used to renderer a reference image are known, and the BRDF available on the user side (target) is different from the source one. Despite of having no access to the implementations of the material

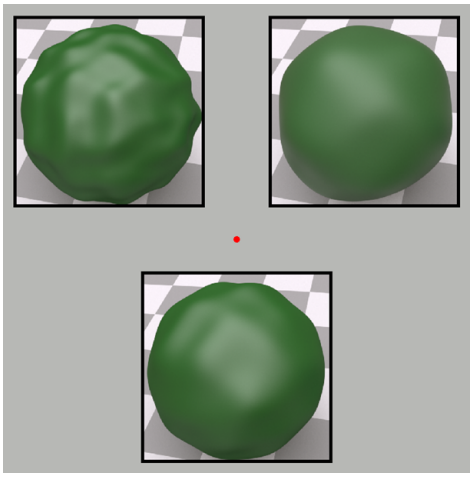

Figure 5.36: Interface of the second psychometric scaling experiment.

models, we proposed an automatic solution to achieve a consistent appearance across different models (and renderers), based on genetic algorithm optimisation; the fitness function used by the genetic algorithm accounts for the differences between source and target BRDFs in the image space.

Unlike recent works, such as [169], requiring an initial guess of the target parameters, our framework provides the user with a tool that can be used either as "black-box" (unconstrained setting) or can incorporate previous knowledge about the semantic and functional relation among source and target parameters, offering additional advantages in the latter case, including a reduced computational cost and the inverse mapping from target to source models.

We tested our framework on several combinations of source/target models, on both dielectric materials and conductors, and perceptually validated our results by means of user studies. In particular these suggest that even in difficult cases, in which the source and target models have extremely different characteristics, such as empirical models as opposed to physically-based models with support to the Fresnel reflectance, our parameters remapping framework allows to obtain a set of parameters for the target model which enable virtually undistinguishable renderings from the ones created with the source model, for a wide portion of the source parameters space. The user studies confirmed the validity of our approach also on complex scenes, closer to a real world scenario.

Our experience with conductors provided some important insight: the differ-

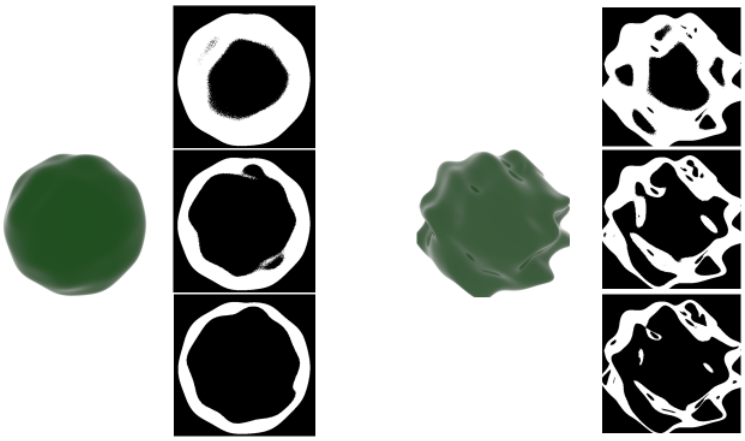

Figure 5.37: Two examples of how each shape is partitioned in two areas by using three different thresholds, given by 102.5%, 110% and 120% of the reflectance at normal incidence (top to bottom). As the threshold increases the "Fresnel" area (in white) moves closer and closer to the silhouette.

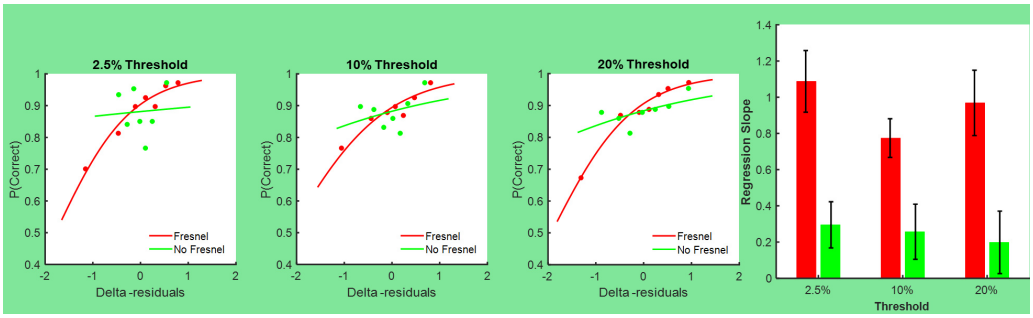


Figure 5.38: Effect of Fresnel on the observers performance in distinguishing between source and remapped renderings.

ences between the statistical distribution of the microfacets used by the source and target model seem to play a fundamental role in the quality achievable by parameters remapping techniques.

# Chapter 6

# Consistent, Tool Independent Virtual Material Appearance

As seen in previous chapters, current materials appearance is mainly tool dependent and requires time, labour and computational cost to deliver consistent visual result despite of the importance in a range of different areas, including Visual Special effects, Interior/Exterior Modelling, Architectural Modelling, Cultural Heritage, Computer Games and Automotive Design, where final decisions about products are often based on a virtual prototype. In this Chapter we address the problem of material consistency across different renderers, in a more general way than in Chapter 5, and describe a solution to achieve it.

## 6.1 Introduction

Currently material appearance heavily depends on the software where virtual material is designed, optimised and rendered (Figure 6.1). A virtual material in most cases does not appear the same as the original once imported in a different renderer due to different algorithms and settings. In fact, each renderer has a specific set of settings and limitations to interpret the properties of a given material data.

Ideally, a standard material model should be consistent and tool independent, to allow coherent material representation. A recent development in this area provides a new file formats for BRDF representation, proposed by NVIDIA®, is the

Material Definition Language (MDL) [170]. MDL is a procedural programming language that allows to define properties of physically plausible material models and integrate them into any supported application. However, it requires some programming skills which may not be part of a typical 3D artist curriculum and, more important, it is not yet widespread. Furthermore, previously rendered scenes cannot benefit from it, for instance in order to achieve consistency throughout the lifecycle of a car which spans several years.

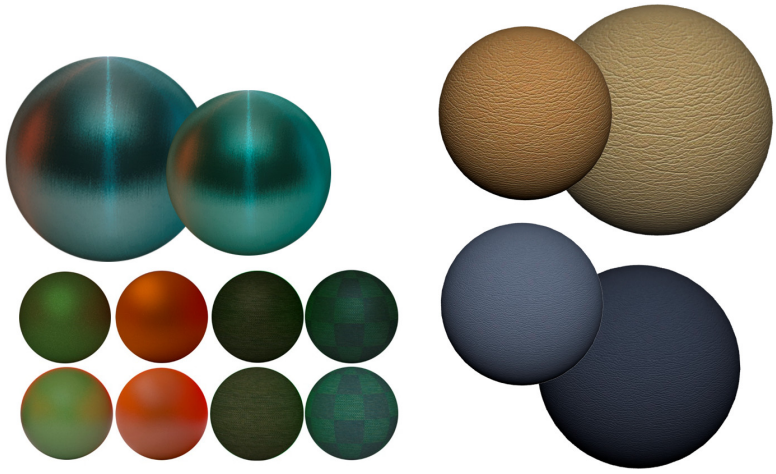

Figure 6.1: Several examples of material rendered in different software.

The aim of this research is to provide artists with a general solution, allowing them to uniform the output of the renderer they use with a "golden standard" application, arbitrarily selected. In this way, within an industry it would be possible to select a reference "golden standard", to which all the renderings obtained with other software will be made visually uniform. In this chapter we propose a solution for the scenario defined in the previous chapter as "maximum-ignorance", summarised as follows:

(i) Knowledge of the source material model and renderer used for the reference rendering;

(ii) No access to the actual implementations of the source nor the target material models.

We propose to characterize the appearance of several classes of materials, rendered using an arbitrary reference software, by extracting relevant visual characteristics. By repeating the same process for any other renderer, ad-hoc mapping functions between the two renderers can be derived. Such mapping functions allow to hallucinate the appearance of a scene, mainly containing the selected classes of materials, under a reference software. My contribution include the following:

- A dataset of reference scenes rendered with different software, carefully selecting the most closely matching settings. Each scene depicts spheres with known materials applied on them, used to characterize the output of each selected renderer with meaningful statistics, per color channel.

- An algorithm aimed to standardise the material appearance. By leveraging existing state-of-the-art algorithms [171], the input rendered image is segmented on a per material basis. With this information, the aforementioned statistics are used to hallucinate the look of each of the materials, providing the user with a post-processed rendering which matches the characteristics of the golden standard renderer in a transparent way.

We address two different real world scenarios, which account for the fact that the reference image might be previously rendered, and hence the information about the environment map and materials used might not be available, and the simpler case of a scene yet to render.

The rest of the chapter is organised as follows. In Section 6.2 we describe some solutions currently available to digital artists trying to achieve consistent material appearance, highlighting their limitations. Section 6.3 describes the structure of a general solution for consistent material appearance, based on the experience gained from Chapter 5 and Section 6.2. Section 6.4 describes our implementation of the general solution, with a detailed report of all the components required. Section 6.5 describes the experiments performed to validate our approach, followed by our conclusions, in Section 6.6.

## 6.2 Assuming the role of a Digital Artist: tentative solutions

In this section we assume the role of an artist, which received an input image from a reference renderer, maybe rendered years before and hence without information about the parameters used to render the materials in the scene, and account for the constraints of this real-world scenario. The most obvious solution, time consuming and error prone, has been described in the previous chapter: it involves the attempt to manually match the appearance using the model available in the rendering tool in use (target), trying to find a new set of parameters to deliver a final appearance as close as possible to the original one.

In order to make the final overall appearance uniform across renderers, several alternative hypothesis have been explored, with an increasing level of complexity a skill required. We start by making use of standard approaches successfully employed, for instance, in the context of a popular commercial graphic editor like Adobe® Photoshop [172]. Among the options it offers there is the possibility to match the appearance between two images, by changing the luminance, changing the color range, and neutralizing a color cast by acting on the global histogram of the input. Unfortunately such solution is bound to fail (see Figure 6.2), since beyond the colour difference, important information such as the location of highlights cannot be modified, the histogram being a global descriptor for the colour content of the image.

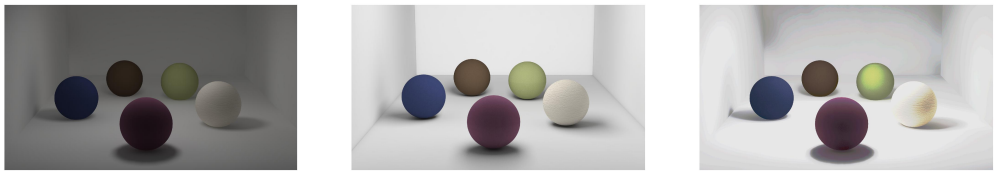

(a) Input Image (b) Reference Image (c) Histogram Matching output

Figure 6.2: Input image (a) and the corresponding output (c) applying a naive histogram matching to the reference image (b).

Hence, on a refined set of experiments the comparison has been based on the color histograms of the different materials in the rendered scenes, assumed to be a

*signature* for each renderer. The key idea is that the pairwise difference between the renderers signatures can be modeled mathematically, hence the inverse formulation will allow to match the appearance of the output of a generic renderer to the reference one, by properly modifying the per material histogram. A similar idea has been exploited previously in different areas of Computer Graphics, however it would be novel in the context of tool-independent representation of physically plausible virtual materials. Several histogram matching algorithms in the state of the art, such as [173] and [174], have been tested for a thorough evaluation. We report here a simple example, in which a set of leather materials were measured and assigned to virtual spheres. The leather samples were rendered in three different applications and, as expected, their visual appearance is different. An example of the same leather material rendered in Autodesk 3D Max, Blender and Autodesk Maya is presented in Figure 6.3.

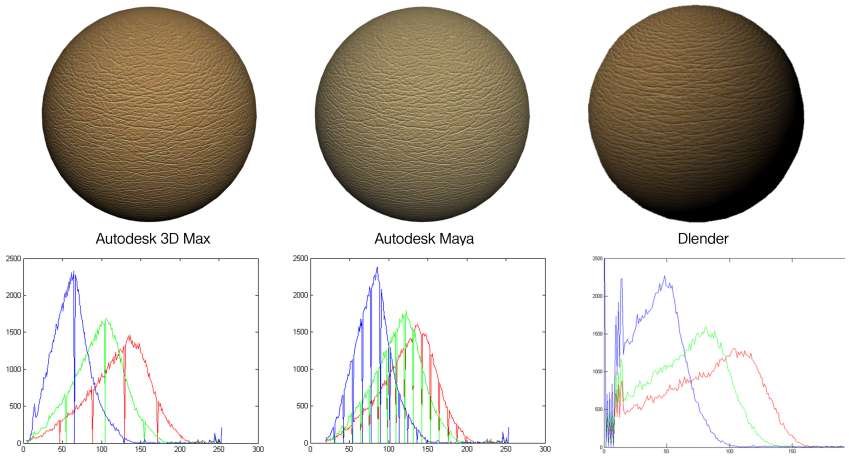


Figure 6.3: Material variation across different renderers. From left to right, the same leather material rendered in Autodesk 3D Max, Blender and Autodesk Maya

We took into account a relatively simple test case, reported in Figure 6.4(a) and (b), where the difference between the two images is rather limited. As expected, this solution does not work to the desired extent. In fact, whereas often it is possible to obtain overall a similar look (Figure 6.4(c)), given by similar global distribution of colors and tones, many characteristics of the BRDF cannot be corrected and made uniform, for instance the reflections over the materials and

(a) Input Image (b) Reference Image (c) Per Material Histogram Matching

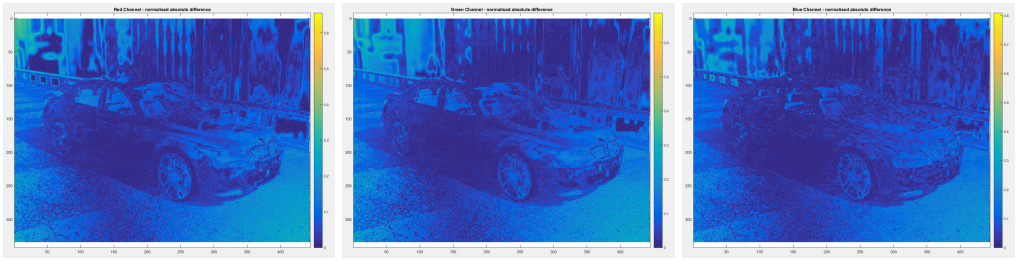

(d) Red Channel - Normalised difference map (e) Green Channel - Normalised difference map (f) Blue Channel - Normalised difference map

Figure 6.4: A simple pair of test images, in which the differences between the output of two renderers is limited. The bottom row shows the normalised colour difference per channel between (b= and (c).

the lighting (Figure 6.4(d-e)). Moreover, in order to work at its best, this naive solution requires the additional key information given by the actual rendering of exactly the same scene with the reference software. Despite of this, a global or even local histogram matching often introduces noticeable artifacts.

A more sophisticated hypothesis was derived from the work by Shih *et al* [157], which triggered the idea to apply a similar technique to scenes rendered with the same set of BRDFs. Given the input image, rendered by the artist in the software in use, the appearance transfer technique could be applied using a database of similar renderings obtained from the reference software. The experiments which have been carried unfortunately did not lead to high quality results, due to the deep differences among the scenes took into account. In fact, the scale of a landscape image is several orders of magnitude larger than a typical automotive image and at the scale considered the effect of the BRDF is much more visible. Hence, physical inaccuracies that do not negatively affect the perception of a good quality at a landscape scale, can have a jarring effect at closer views.

## 6.3 Motivations and Characteristics of a General Solution

The lesson learned from previous section is that in order to achieve virtual material appearance consistency a general solution must be formal, based on mathematical framework with at its core systematic measurements of differences among renderers, rather than ad hoc examples. Based on the experience acquired with our BRDF parameters remapping work (chapter 5), a different framework has been hypothesised. We focused on the following components:

- A set of reference materials;
- A set of renderings of the reference materials, using the most common renderers;
- Suitable metrics to analytically characterise the output of a given renderer and the difference among different renderers output;
- Derive an analytical formula to match the appearance an arbitrary reference renderer.

We propose to characterize the appearance of several classes of materials rendered using an arbitrary reference software by extracting relevant visual characteristics. By repeating the same process for any other renderer ad-hoc mapping functions between the two renderers can be derived. Such framework allows to hallucinate the appearance of a scene, depicting mainly the selected classes of materials, under a reference software (Figure 6.5).

## 6.4 Proposed solution

We make the following hypothesis: a rendering system, consisting of a set of BRDF models with unknown implementation, is assumed to act as a set of filters, on a per material basis, on the rendered image. Hence, if the same scene is rendered using two different tools with settings as close as possible, the visual differences can be attributed to the effects of the aforementioned filters, which we aim

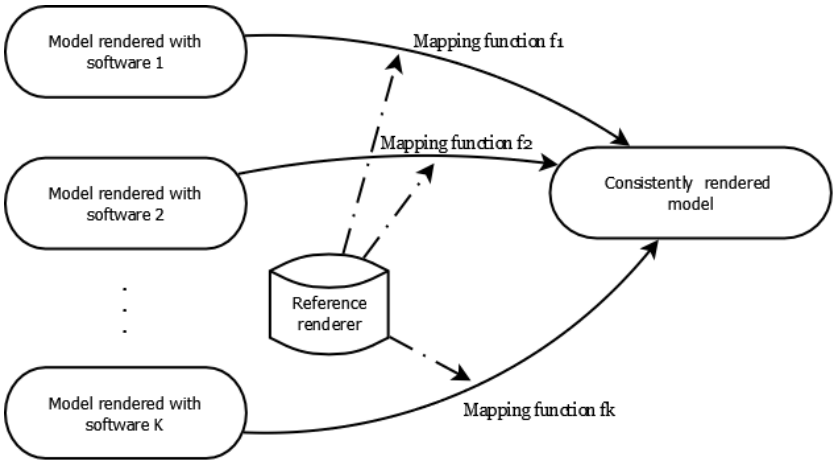


Figure 6.5: Diagram of a general framework for consistent material appearance across different renderers.

to estimate. Given any two renderers $R_a$ and $R_b$, set of materials $M_i$, $i = 1,\ldots,k$, the knowledge of such filters $f_{a,M_i}$ and $f_{b,M_i}$ allows to remove the effect of $R_a$ from an image $I_a$, by computing $\hat{I}_a = {f_{a,M_i}}^{-1}(I_a)$, $i = 1,\ldots,k$; the subsequent application of $f_{b,M_i}$ on $\hat{I}_a$ provides an image which visually mimics the typical output of $R_b$.

In the following Sections the main components and steps of the proposed approach will be described; such steps are summarized in the following bullet points:

- Characterisation of the renderer $R$ in terms of the parameters of the associated filter $f_R$. The lighting and the geometry used for the characterisation are known. This process is performed only once for each renderer.

- Selection of a reference and a target renderers;

- Material segmentation on the input image, rendered with the target renderer, which needs to be made visually consistent with the reference one, in order to apply a per material correction;

- Lighting estimation from the input image; the materials belonging to the car are removed from the scene;

- Rendering of a set of spheres, one per each material in the scene, using the statistics of the selected reference and target renderers and the estimated

lighting;

- Appearance transfer at a pixel level. The surface normal $\mathbf{n}_p$ at pixel $p$ and its material ID $m_p$ are used as entries for lookup table approach on the rendered sphere with the same ID. This process is repeated for each material in the scene, leading to the hallucination of the scene as rendered by the reference renderer.

The diagram of the proposed framework is reported in Figure 6.6; the blue side refers to the characterisation part, the green side to the appearance hallucination part for previously rendered images, which provides a consistent appearance across renderers. A more detailed diagram for this part is reported in Figure 6.7, where also the case of appearance hallucination of a scene yet to render is reported. Algorithms 2 reports the pseudo-code for the characterization part; Algorithm 3 and 4 respectively report the case of an image previously rendered, possibly with missing information (lighting, materials, *etc.*), and the case of an image non yet rendered, with all the information available.

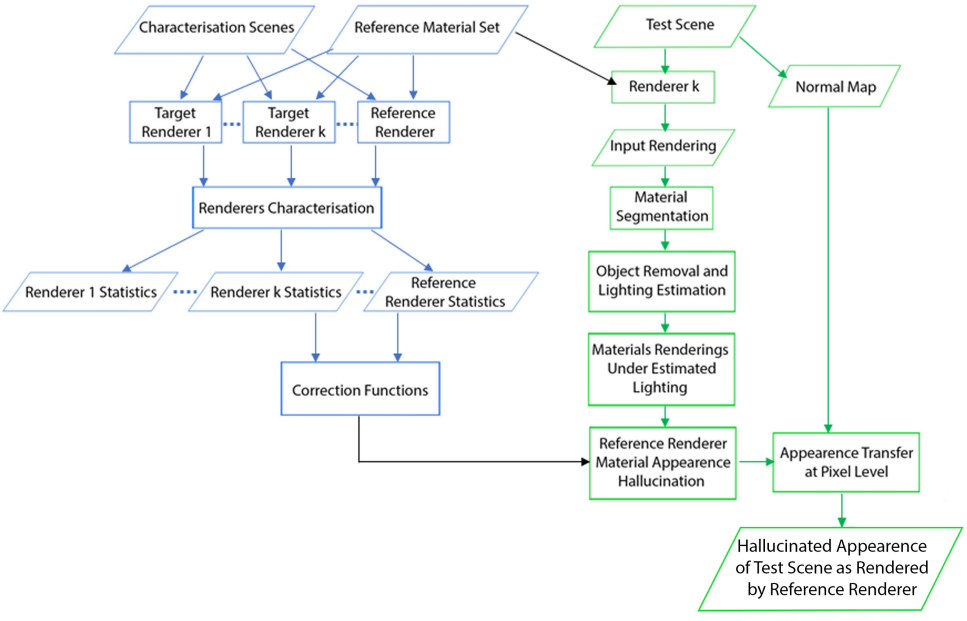


Figure 6.6: Diagram of the proposed framework for consistent material appearance across different renderers.

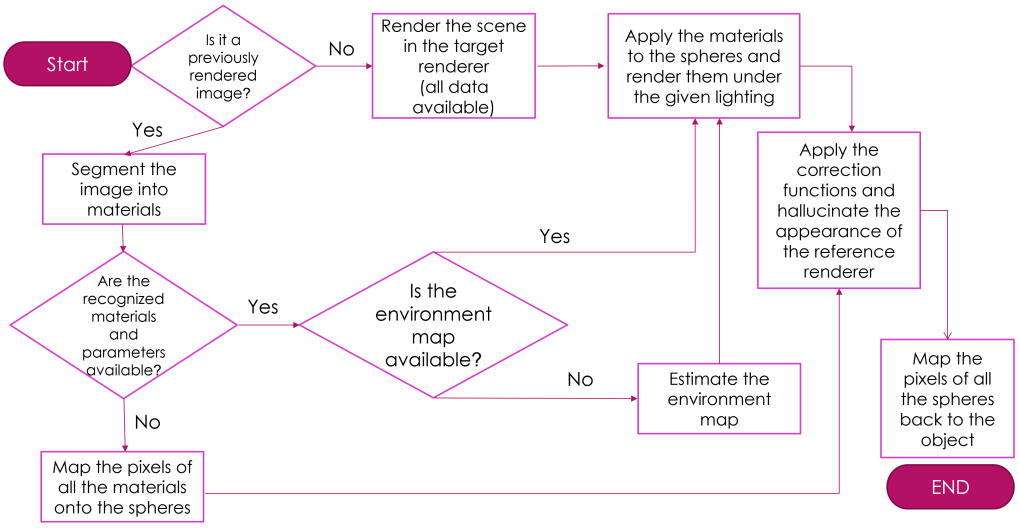


Figure 6.7: Diagram of the proposed framework application, taking in to account all possible scenarios, for consistent material appearance across different renderers (Chapter 6).

**Algorithm 2** Pseudo-code of the characterization part of our framework for consistent virtual material appearance. Input to the algorithm are a scene geometry (*sceneGeom*, including shapes, positions, lighting), a set of $|K|$ materials ($\{MaterialsSet\}$)and a set of $|I|$ renderers ($\{RenderersSet\}$), one of which is the reference renderer (*refRend*); in the following, sets are indicated by $\{\cdot\}$. The output of the algorithm is a set of $|I \times K|$ correction functions, one for each renderer and material, used to hallucinate the appearance of the corresponding material in the reference renderer.

**for** each material $M_k \in \{MaterialsSet\}$ **do**
  $renderedMaterialImage_k \leftarrow renderImage(sceneGeom, M_k, refRend)$
  ${\alpha^*}_k \leftarrow computeEnergyStat(renderedMaterialImage_k)$
  ${\lambda^*}_k \leftarrow computeMeanStat(renderedMaterialImage_k)$
  ${\psi^*}_k \leftarrow computeVarianceStat(renderedMaterialImage_k)$
**end for**
**for** each renderer $R_i \in \{RenderersSet\}$ **do**
  **for** each material $M_k \in \{MaterialsSet\}$ **do**
    $renderedMaterialImage_{i,k} \leftarrow renderImage(sceneGeom, M_k, R_i)$
    $\alpha_{i,k} \leftarrow computeEnergyStat(renderedMaterialImage_{i,k})$
    $\lambda_{i,k} \leftarrow computeMeanStat(renderedMaterialImage_{i,k})$
    $\psi_{i,k} \leftarrow computeVarianceStat(renderedMaterialImage_{i,k})$
    $f_{\alpha i,k} \leftarrow computeCorrectionFunction(\alpha_{i,k}, {\alpha^*}_k)$
    $f_{\lambda i,k} \leftarrow computeCorrectionFunction(\lambda_{i,k}, {\lambda^*}_k)$
    $f_{\psi i,k} \leftarrow computeCorrectionFunction(\psi_{i,k}, {\psi^*}_k)$
  **end for**
**end for**

**Algorithm 3** Pseudo-code of the application part of our framework for consistent virtual material appearance, addressing the case of an image previously rendered, possibly with missing information. Input to the algorithm are an image ($refImage$), rendered with the render currently in use ($R_i$), the correction functions derived by means of Algorithm 2 and the input image normal map ($normMap$). The input image needs to be made consistent with the appearance which would be produced by reference renderer ($refRend$, see Algorithm 2). Optional parameters are a set of materials parameters used for the rendering ($\{MaterialParams\}$, possibly empty) and the environment map used in the rendering ($envMap$, possibly unknown). In the following, sets are indicated by $\{\cdot\}$; materials $M_k$ are considered to be objects with two fields, $M_k.params$, containing the parameters of the material (if known) and $M_k.locs$, containing the locations of all the pixels in $refImage$ which depict the material $M_k$. The output of the algorithm is $outputRendering$.

$outputRendering \leftarrow refImage$
$refImageCopy \leftarrow refImage$
$\{MaterialsSet\} \leftarrow segment(refImage)$
**if** $\nexists envMap$ **then**
  **for** each material $M_k \in \{MaterialsSet\}$ **do**
    $refImageCopy \leftarrow removeMaterial(refImageCopy, M_k.locations)$
  **end for**
  $envMap \leftarrow estimateEnvMap(refImageCopy)$
**end if**
**for** each material $M_k \in \{MaterialsSet\}$ **do**
  **if** $\exists \{MaterialParams\} and M_k.params \in \{MaterialParams\}$ **then**
    $sceneSphere_k \leftarrow NULL$
    $sceneSphere_k \leftarrow apply(M_k.params, sceneSphere_k)$
    $imageSphere_k \leftarrow render(sceneSphere_k, R_i, envMap)$
  **else**
    $imageSphere_K \leftarrow map(M_k.locs, refImage, normMap)$
  **end if**
  $imageSphere_k \leftarrow applyCorrectionFunctions(f_{\alpha i,k}, f_{\lambda i,k}, f_{\psi_{i,k}}, imageSphere_k)$
  $outputRendering \leftarrow inverseMapping(M_k.locs, imageSphere_k, normMap)$
**end for**

**Algorithm 4** Pseudo-code of the application part of our framework for consistent virtual material appearance, addressing the case of an image not yet rendered, with all the information available (scene geometry, lighting, material parameters). Input to the algorithm are an image (*refImage*), rendered with the render currently in use ($R_i$), the correction functions derived by means of Algorithm 2, the input image normal map (*normMap*), the set of materials used for the rendering and their parameters ($\{MaterialsSet\}$) and the environment map used in the rendering (*envMap*). The input image needs to be made consistent with the appearance which would be produced by reference renderer (*refRend*, see Algorithm 2). The notation used and the output are the same as in Algorithm 3.

$outputRendering \leftarrow refImage$
**for** each material $M_k \in \{MaterialsSet\}$ **do**
  $sceneSphere_k \leftarrow NULL$
  $sceneSphere_k \leftarrow apply(M_k.params, sceneSphere_k)$
  $imageSphere_k \leftarrow render(sceneSphere_k, R_i, envMap)$
  $imageSphere_k \leftarrow applyCorrectionFunctions(f_{\alpha\,i,k}, f_{\lambda\,i,k}, f_{\psi_{i,k}}, imageSphere_k)$
  $outputRendering \leftarrow inverseMapping(M_k.locs, imageSphere_k, normMap)$
**end for**

### 6.4.1 Statistical Analysis of Renderers

Zubiaga *et al.* [154] focused on how the properties of BRDFs influence the rendered picture, by working locally in Fourier space and analysing how BRDF moments up to order 2 induce colouring, warping and blurring of reflected radiance. In their work a subset of unimodal materials in the MERL database [4] has been used, limiting the analysis to 2D slices of the selected BRDFs, assumed to act as filters on the incident lighting, whose parameters need to be estimated. Although a BRDF is a 4D function, the choice of 2D slices of the BRDF is justified by the consideration that when the radiance reaching the eye from a surface point is computed, the view direction is kept fixed. The 2D slices of the selected BRDFs are pre-processed with a heuristic method for diffuse and specular separation, and parameterised using a view-centred angular parameterisation with poles orthogonal to the view direction, to minimise distortions around the scattering plane.

Building upon [154], Zubiaga *et al.* [153] noted that the image $I$ of a sphere of a material $M$, in orthographic projection, rendered under some environment lighting $L$, can be written as a 2D spherical convolution under the radial symme-

try hypothesis: $I = M * L$. Thanks to the radial symmetry it is possible to restrict the analysis to a $1D$ slice of $M$; in particular, if an angular parametrization $(\theta, \phi)$ based on screen-space normals is used, the analysis can be limited to the $\theta$ direction.

My research aims to understand how a specific rendering software influences the final image, hence I made the hypothesis that the image $I_R$, obtained with the rendering tool $R$, can be written as:

$$I_R = f_R * M * L. \tag{6.4.1}$$

where $f_R$ is the set of filters (one per material) introduced by $R$.

By keeping the material properties as coherent as possible across renderers, and thanks to the known lighting, the statistic proposed in [154] to characterize $M$ can be readily adapted to define the properties of $R$ instead, on which we focus (figure 6.8). In particular, the energy $\Gamma$, mean $\mu$ and variance $\rho$ for each point of the sphere, can be derived as described in [153, 154]:

$$\Gamma(I_R) = \Gamma(L_\phi)\alpha(\theta), \tag{6.4.2}$$

$$\mu(\bar{I}_R) = \mu(\bar{L}_\phi) - \lambda(\theta), \tag{6.4.3}$$

$$\sigma(\bar{I}_R) = \sigma(\bar{L}_\phi) + \psi(\theta), \tag{6.4.4}$$

where $L_\phi$ is the lighting $L$ integrated over the $\phi$ direction, $\bar{L}_\phi$ and $\bar{I}_R$ are normalized by the energy $\Gamma(I_R)$; the parameters of the filter $f_R$ are given by $(\alpha, \lambda, \psi)$, which are related to important features like coloring, warping and blurring of the reflected radiance. The statistical analysis is performed on both the diffuse and specular component separately, which on our dataset (described in Section 6.4.2) are both known.

### 6.4.2 Training Set

In order to derive the filters characteristics, a dataset tailored for the automotive industry has been built as a case study. Six material classes have been selected, namely car paint, plastic, wood, glass, fabric and leather, since they are the most

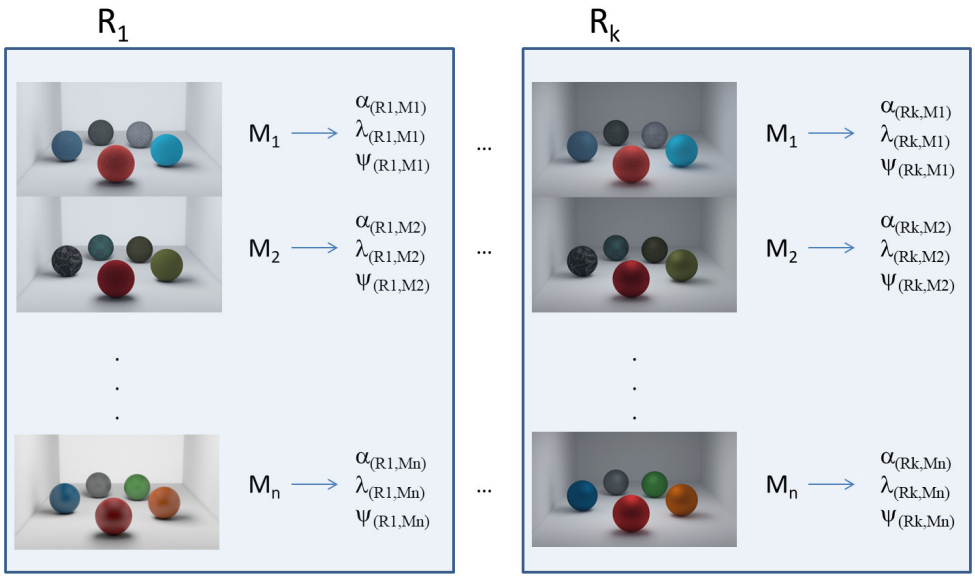


Figure 6.8: Renderer characterisation. Our dataset, containing 5 variants per each class of materials, is used to characterize a given renderer.

common ones in automotive industry; within each class we consider 5 variants, which differ in terms of diffuse and specular components.

Each material is turn assigned to a sphere and placed at the center of the scene, which consists of a non-standard Cornell box where all walls are set, in turn, to the same primary RGB color; additionally, a setting with the sphere only is included, illuminated by the environment map reported in Figure 5.2. The rational is that such illumination condition carry important information about the BRDF spectral response, which can be analysed separately for each color channel, thanks to the Cornell boxes, and all in one image thanks to the environment map described in chapter 5, successfully used in the context of our BRDF remapping framework.

Hence, in total we have 6 classes, 5 samples per each class and 4 settings for the lighting, with a total of 120 scenes, each of which rendered with all the selected renderers, chosen among the most used ones: Autodex 3Ds Max, Autodex Maya and Blender (see figures 6.9, 6.10, 6.11 and 6.12 for same examples and visual comparisons).

Our dataset allows us to characterize, for each class of materials, the influence of a renderer on the visual appearance, along with simple information such as global and local histograms per color channel, aimed to address the rest of

materials in a scene which are not included in our set.

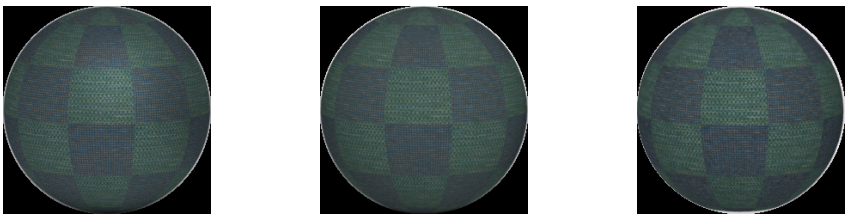

(a) Blender (b) Maya (c) 3ds Max

Figure 6.9: From left to right: fabric samples inside the white box rendered respectively with Blender, Autodesk Maya, 3ds Max. All the renderings have been linearly scaled in order to have the same average value per each color channel.

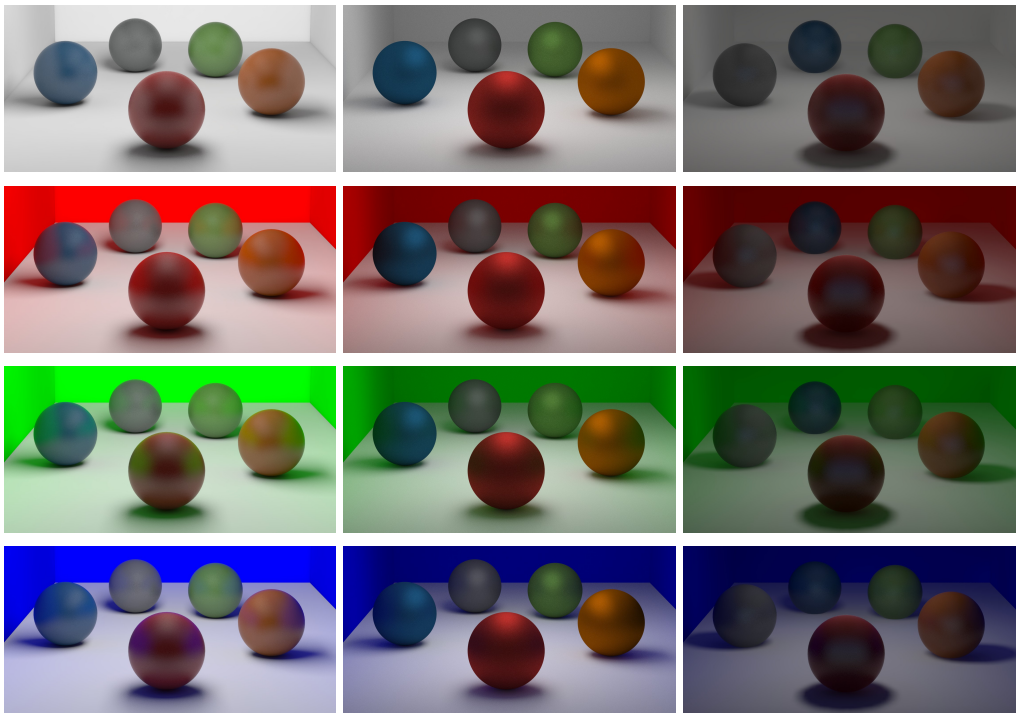

Figure 6.10: The 5 spherical samples for the category “car paint”, inside non-standard Cornell boxes. The images in the first column refers to Autodesk Maya, in the second column to Blender, in the third column to 3ds Max.

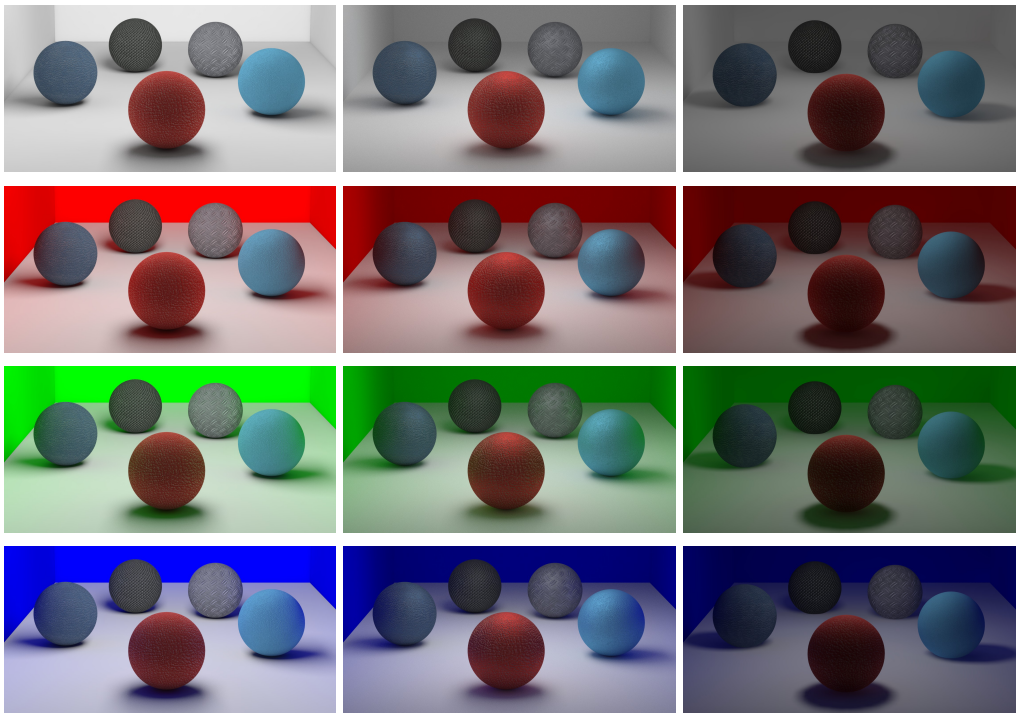

Figure 6.11: The 5 spherical samples for the category “plastic”, inside non-standard Cornell boxes. The images in the first column refers to Autodesk Maya, in the second column to Blender, in the third column to 3ds Max.

Textured materials with a rough surface, like the fabric sample in Figure 6.9, display mesoscopic effects like inter-reflections, self-occlusions and self-masking, hence a Bidirectional Texture Function (BTF) should be used to describe their properties, thus preventing a direct application of the framework described in [154]. However, as observed by Dana *et. al.* [132] the BRDF is able to describe material variation of a textured material at a coarse scale, hence averaging through the sample leads to its BRDF.

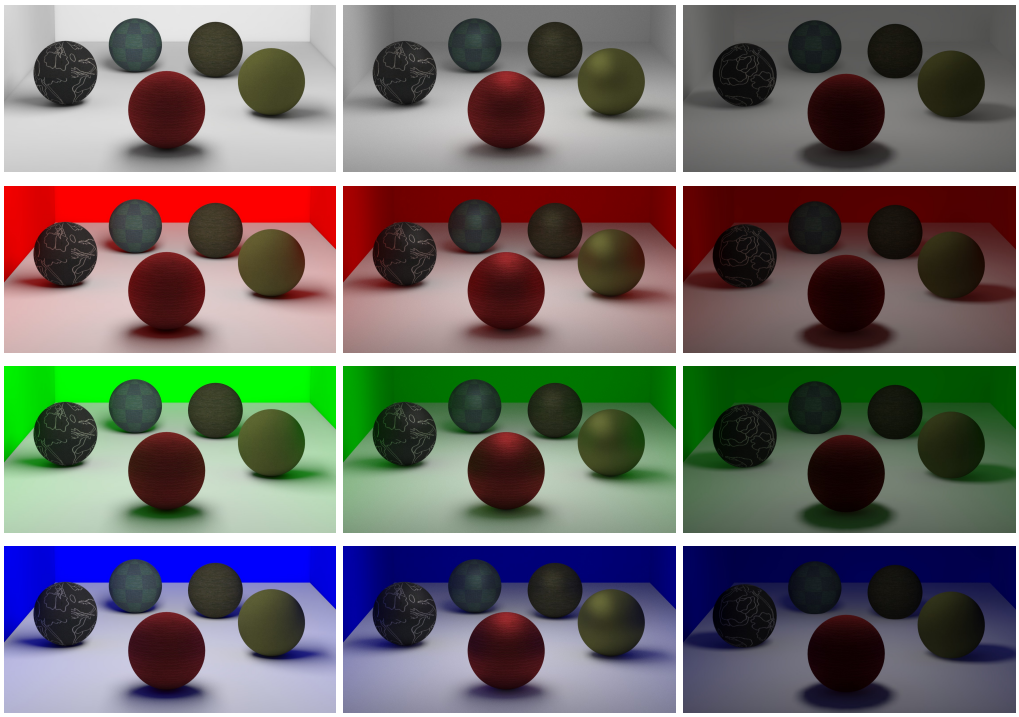

Figure 6.12: The 5 spherical samples for the category "fabric", inside non-standard Cornell boxes. The images in the first column refers to Autodesk Maya, in the second column to Blender, in the third column to 3ds Max.

### 6.4.3 Image Segmentation and Materials Map

In previous Sections we have described how the influence of a renderer on the visual appearance is characterized, *i.e.* on a per material basis. Accordingly, the derived correction functions need to be applied on a per material basis, as reported in Algorithms 3 and 4. In order to achieve this, it is required to know where the materials we account for are located in the input image. For this purpose, a segmentation algorithm addressing the materials depicted in the input needs to be devised, in order to group together all the pixels belonging to the same material class.

Both the works by Zubiaga *et al.* [153] and Khan *et al.* [6] require either a map with material IDs or a foreground/background map, which could be challenging to create manually for a complex scene, if not provided by the rendering tool in use. However, powerful computer vision techniques have recently become available to

help in this task. In the last few years, several dataset of texture images, including man-made materials, have been released.

The Material in Context Database (MINC) is a large dataset containing about 3 millions material samples, used to train a convolutional neural network (CNN) and a conditional random field (CRF), combined together to recognize and segment materials in the wild [175]. In particular, a CNN is trained in order to produce a single prediction for a given input patch. The trained CNN is used as a sliding window to predict material across the image, at different scales. The prediction across the scales are then averaged and given in input to the CRF, in which all pairs of variables are directly connected by pairwise potentials; the fully connected pairwise reasoning outputs material predictions for each pixel.

The material categories considered in this study are sufficient to faithfully describe most of the appearance of a car model, and at the same time they are included in the MINC dataset [175], thus enabling a straightforward use of the GoogLeNet [176] network for the image segmentation into material classes, useful to obtain a map with material IDs from an unlabelled input rendering. We augmented the MINC dataset with 100 images including car paint and trained the segmentation network.

Once the material labels are derived by means of the segmentation algorithm (Figure 6.13), the relevant material pixels can be removed from the image to estimate the lighting (Section 6.4.4 and mapped on a set of spheres (Section 6.4.5), one per material, in order to apply select the correct statistics for each of them (Figure 6.13).

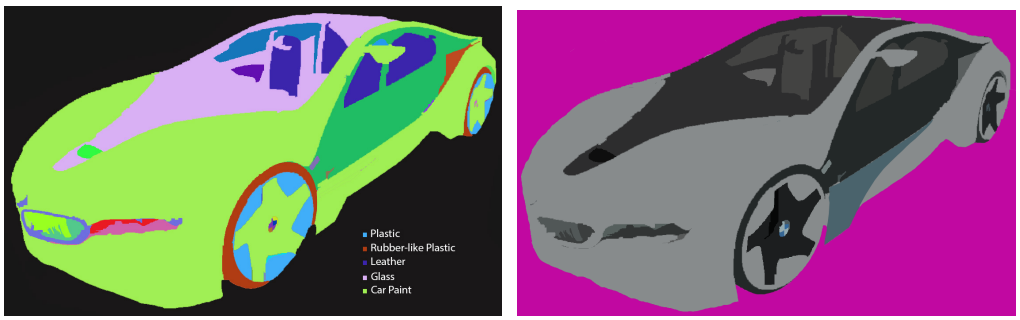


(a) (b)

Figure 6.13: Materials ID map (a) and segmented image (b).

### 6.4.4 Incident Lighting Estimation

In order to estimated the incident lighting, we follow the approach described in [6]. All the background pixels together provide direct information for a subset of incident direction. The pixel belonging to the object and the area outside the image prevent a complete estimation of the incident lighting, which hence needs to be approximated. The object is removed from the image $I$ thanks to the information deriving by the material segmentation, and the hole left by the removal process is filled by a simple inpainting technique, which copies pixels both from the left and from the right parts of the hole, using blending weights determined by the distance from the original location in the background to the new position in the hole [6]:

$$I_I(x,y) = w_1 I((2x_m - x) \mod X, y) + w_2 I((2x_M - x) \mod X, y) \qquad (6.4.5)$$

where $I$ and $I_I$ have size $X \times Y$, $x_m$ is the horizontal coordinate of the leftmost point in the hole, $x_M$ the coordinate of the rightmost point, $w_1 = \frac{x - x_m}{x_M - x_m}$, $w_1 = \frac{x_M - x}{x_M - x_m}$ and the mod operator allows to deal with coordinates outside the bounds, hence treating the image as periodic. Despite of the simplicity, this method performs well when the size of the object removed is relatively small compared to the dimensions of the image, and the content of the image is visually rich; a simple environment in which there is a limited number of colours and sharp edges are dominant can introduce some periodic pattern (see Figure 6.14).

From the inpainted image $I_I$, placed on the $xy$ plane, the biggest possible circle is cut from its central area (*i.e.* the center of the image and the center of the circle coincide and the diameter of the circle is equal to the smallest dimension of the image) and it is mapped onto a hemisphere, extruded along the $z$-direction (Figure 6.15); the conventional 3D sphere used for image based lighting is obtained by mirroring the hemisphere along the $z$-axis.

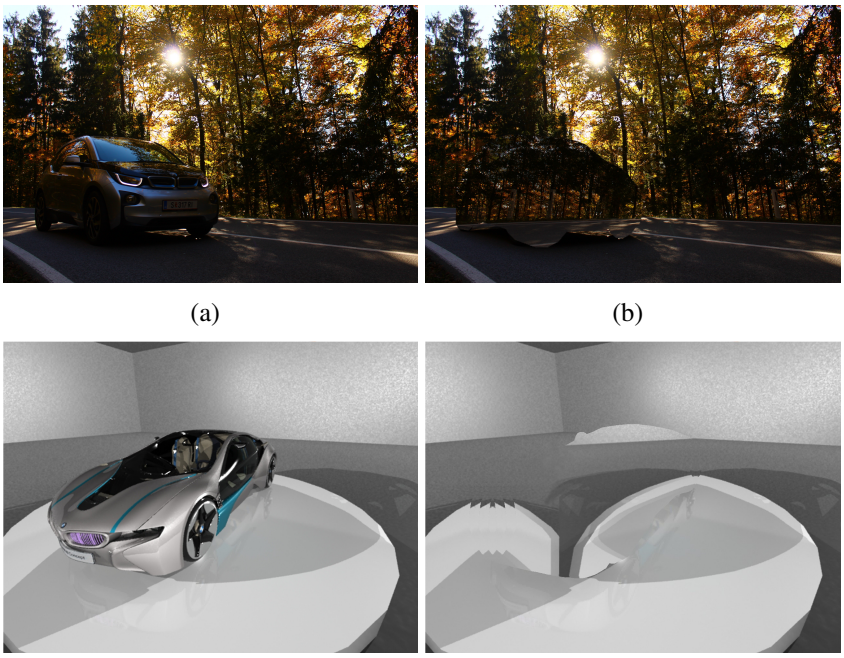

(a) (b)

(c) (d)

Figure 6.14: A simple inpainting technique to estimate the light form behind the object [6]. On the left side and example in which the object is placed in a visually rich image ((a) photo by Maximilian Wachter on Unsplash.com); on the right an example in which the characteristics of the environment introduce artifacts, such as periodic patterns.

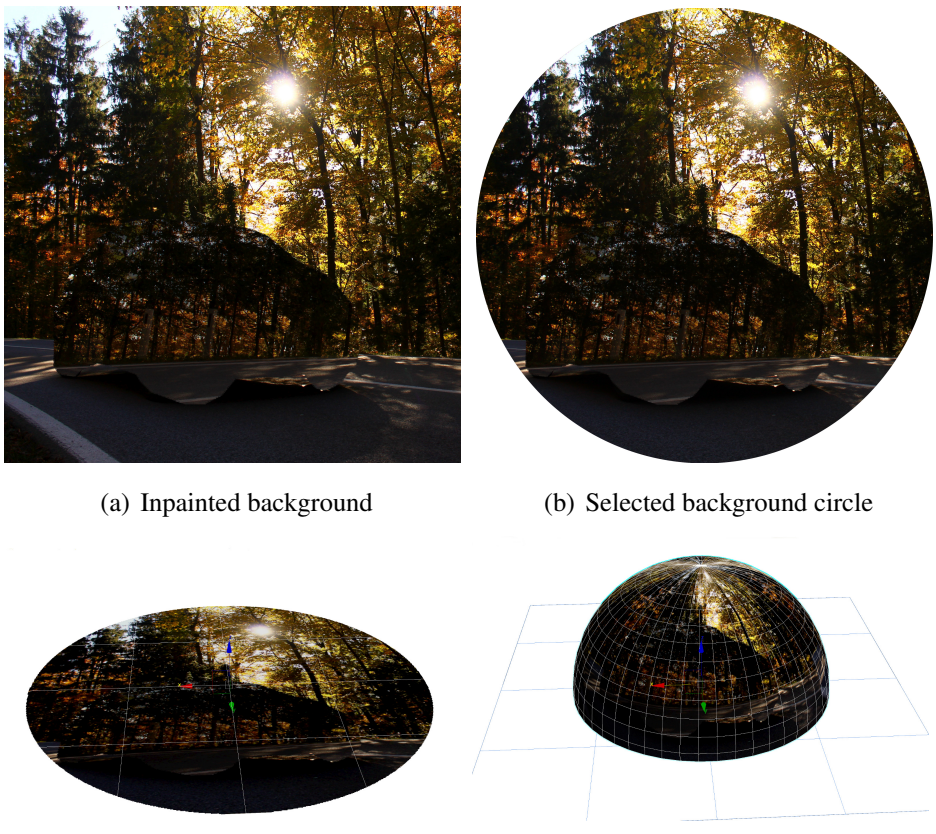

(a) Inpainted background (b) Selected background circle

(c) Selected circle on the plane (d) Selected circle mapped on a hemisphere

Figure 6.15: The inpainted background is used to create an approximated environment map. A circle with the radius equal to the height of the image is selected from around the object removed, placed on plane, and then mapped on a hemisphere, which is mirrored to create the final environment map, as described in [6].

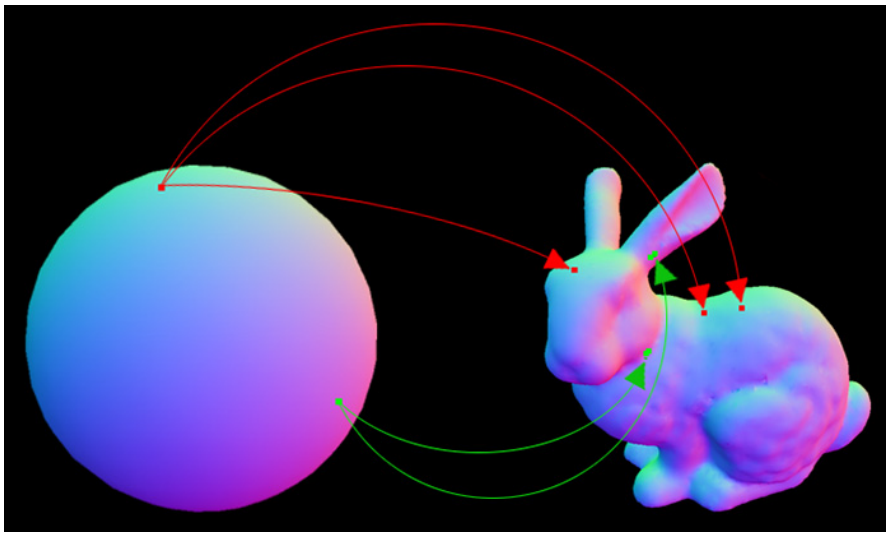

(a) Normal-based lookup

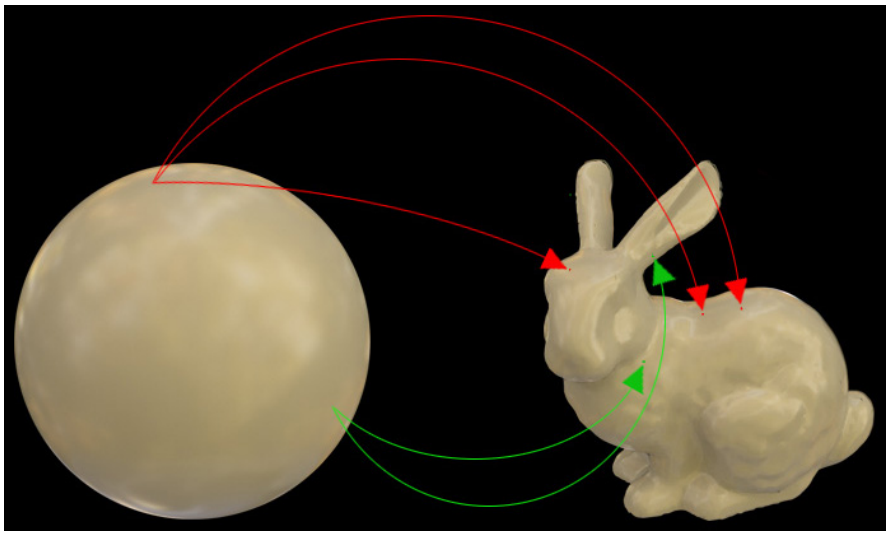

(b) Appearance transfer

Figure 6.16: The surface normals are used to map the appearance of the sphere of the ball under the estimated lighting on the object; this process is analogous to a MatCap assignment (Section 4.1). The surface normals are encoded as $\frac{1}{2}\mathbf{x}+1 \rightarrow Red$, $\frac{1}{2}\mathbf{y}+1 \rightarrow Green$, $\frac{1}{2}\mathbf{z}+1 \rightarrow Blue$. The plastic sphere in the bottom row is rendered using the environment map reported in Figure 6.15.

### 6.4.5 Appearance transfer at pixel level

At this stage, the only missing information is the diffuse and specular components of the materials in the input rendered image. In this work the focus is not on

diffuse/specular separation but rather on providing the user with a tool to select the diffuse color for each material, corresponding to $\alpha(0)$ in 6.4.2, subtracted from the corresponding pixels and hence providing an approximation of the specular component.

For each pixel $p_o$ belonging to the object for which the material appearance needs to be hallucinated, from the previous stages of the pipeline we know the material IDs and surface normals $\mathbf{n}_{p_o}$, which allow us to perform a mapping $(n_{p_o,x}, n_{p_o,y}, n_{p_o,z}) \rightarrow (n_{p_s,x}, n_{p_s,y}, n_{p_s,z})$ on the sphere normals $\mathbf{n}_{p_s}$(Figure 6.16(a)) and generate a corrected value by applying Equations 6.4.2, 6.4.3 and 6.4.4 for the reference render, accounting for the estimated incident lighting as in Equation 6.4.1.

The new pixel value is then applied to the input rendering image (Figure 6.16(b)), thus hallucinating the appearance of a typical output of the reference renderer.

## 6.5 Experiments and Results

Our method has been tested using typical scenes from the automotive industry. Our preliminary results show the effectiveness of the proposed method. In Figure 6.17 a sketch car interior rendering 6.17(a) has been processed and the hallucinated output 6.17(d) can be visually compared to the ground truth obtained from the reference render (Figure 6.17(b)); in the same figure it is possible to see the output obtained by inverting the roles of the renderers (Figure 6.17(c)). In Figure 6.18 a different car model and scene are reported; the error maps before and after applying our method are reported in Figure 6.19.

The proposed solution can be readily extended to other renderers and the reference software can be seamlessly replaced by a different one. Although our approach provides promising preliminary results, the employed approximations inevitably lead to some limitations, which we aim to address in future work.

A first limitation is clearly due to the set of materials considered in this study, tailored for the automotive industry. A more general solution would require a broader set of materials, which could potentially lead to a vast amount of raw data to analyse and store.

The inclusion of other classes of materials could however pose additional is-

sues related to the material segmentation, in particular in case of materials not included in the MINC dataset: in these cases, it is important to keep the dataset well balanced not to affect the segmentation performance, thus requiring thousands of samples to label; we partly faced this issue when including the "car paint" class, which can be misclassified (*e.g.* the assigned material label could be "metal", more represented in the dataset).

Another source of inaccuracy is due to the simple estimation of the environment lighting, which also poses the additional constraint of having a background area considerably bigger than the object to inpaint: such a limitation is particularly relevant in case of close up scenes.

Similarly to [153], my approach cannot mimic inter-reflections or shadowing effects, thus leading to noticeable errors in large shadow areas (see Figure 6.17). This could be mitigated by using an artist-created occlusion map. The current implementation requires a normal map in input, which can prevent the application of the proposed technique in case of already existing rendered image for which it is only known the source renderer. In such situations a shape from shading approach could allow to derive a depth map and from this an approximated normal map [6], although it could be particularly challenging in case of metals and glass, a particularly common situation in the automotive industry.

## 6.6 Conclusion

Please note that the described method does not require to render the scene with the reference software, which is only used to derive the dataset for characterization and to obtain the ground truth in the experiments. Even with the additional key information given by the actual rendering with the reference software, a naive solution, such as a simple global (or even local) histogram matching, clearly cannot deliver accurate material appearance and introduces noticeable artifacts (see Figure 6.2).

A possible limitation is that each new release of a rendering tool might require a new characterization, although each new candidate release generally involve the use of benchmarks to guarantee visual consistency with previous versions. Even in case of significative differences between different releases, this would simply

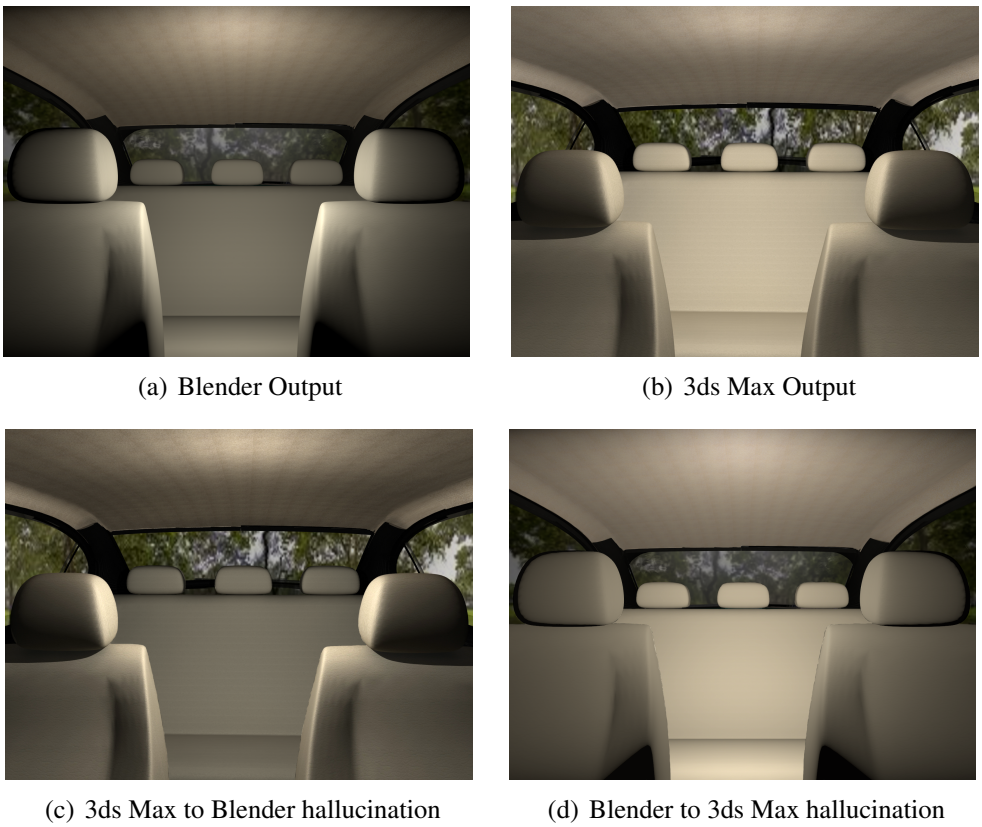

(a) Blender Output (b) 3ds Max Output

(c) 3ds Max to Blender hallucination (d) Blender to 3ds Max hallucination

Figure 6.17: Car interior scene from the design stage. In (a) the scene is rendered with Blender, in (b) with 3ds Max; In (c) the image (b) has been processed in order to match (a). In (d) the image (a) has been processed in order to match (b).

require a renderer/release characterization, which could have the only side effect of increasing the size of the dataset and of the database of parameters to store.

Up to now there are only few methods and tools that allows to edit BRDFs, typically in an interactive way on 3D scenes, as reported in [5]. The proposed method does not edit BRDF materials but rather hallucinate their appearance in a post-rendering step, making them visually consistent with a reference software. Despite of the current limitations, the proposed work is a significant step towards a tool independent, consistent material appearance across different renderers, in an intuitive and transparent way for artists.

An interesting application for our “maximum ignorance” framework would be the use also in the “minimum ignorance” scenario described in chapter 5, thus

hallucinating the appearance of the source BRDF rather than remapping the parameters.

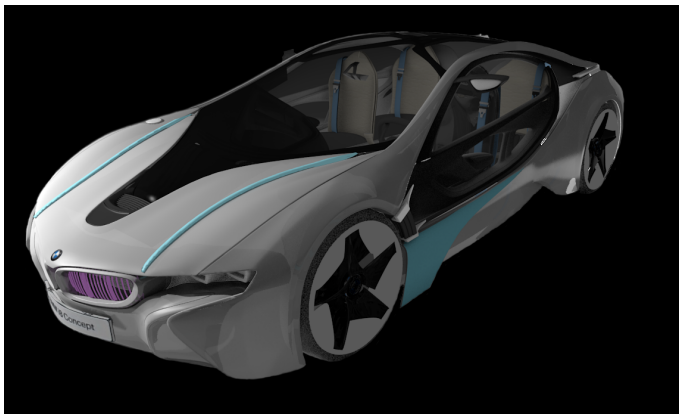

(a) Blender Output

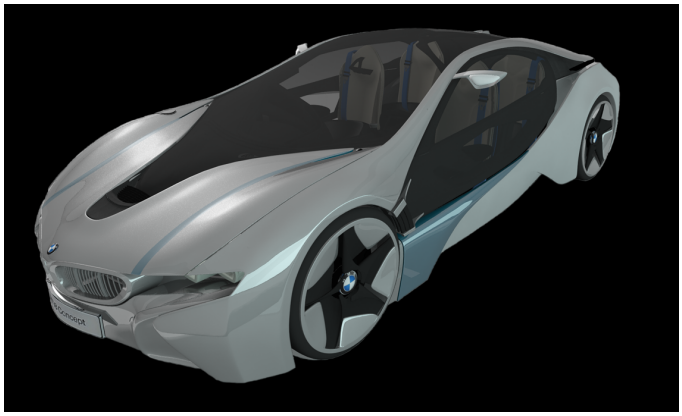

(b) 3ds Max Output

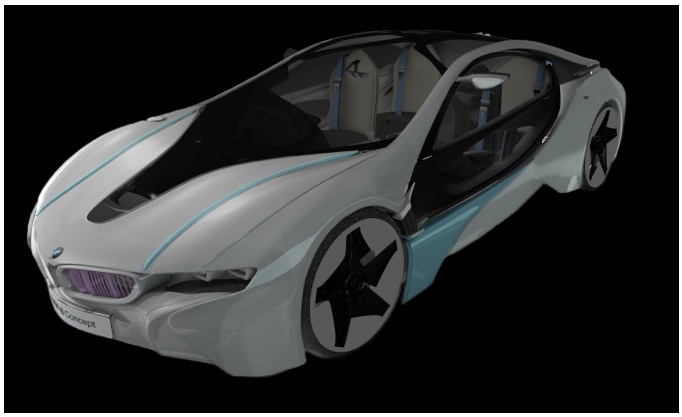

(c) Blender to Max hallucination

Figure 6.18: Car exterior scene in an advanced design stage. In (a) the scene is rendered with Blender, in (b) with 3ds Max; In (c) the image (b) has been processed in order to match (a). Please note that the view point in (a) and (b) are slightly different, as well as the lighting. The image has been cropped for visualization purposes, removing part of the background.

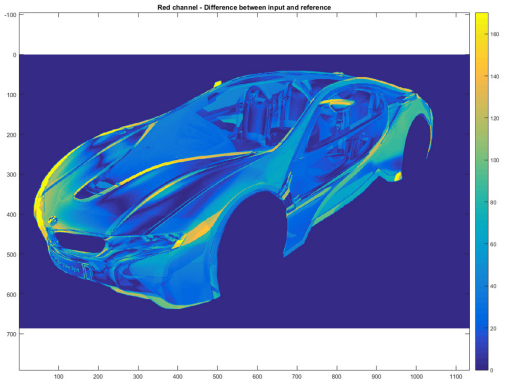


(a) Input - NRMSD Red Channel 0.1196

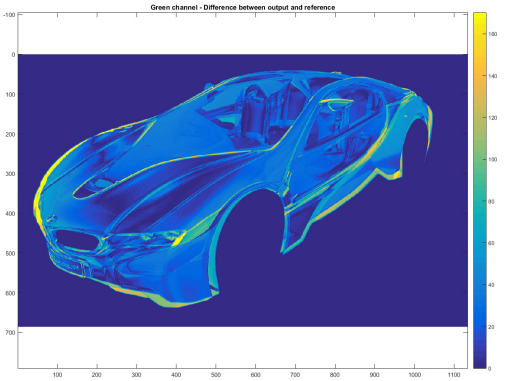


(b) Output - NRMSD Red Channel 0.0978

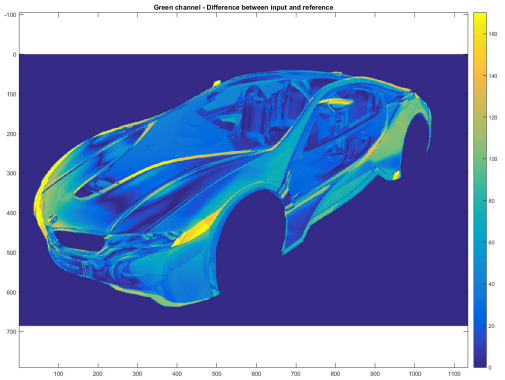


(c) Input - NRMSD Green Channel 0.1268

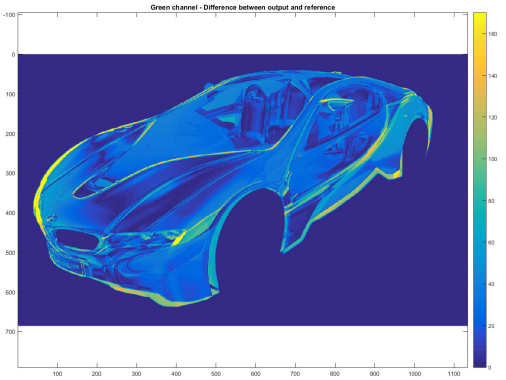


(d) Output - NRMSD Green Channel 0.1034

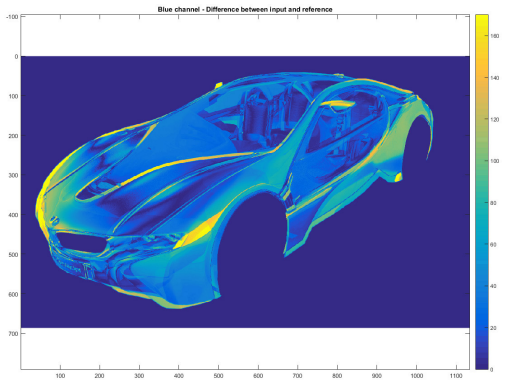


(e) Input - NRMSD Blue Channel 0.1273

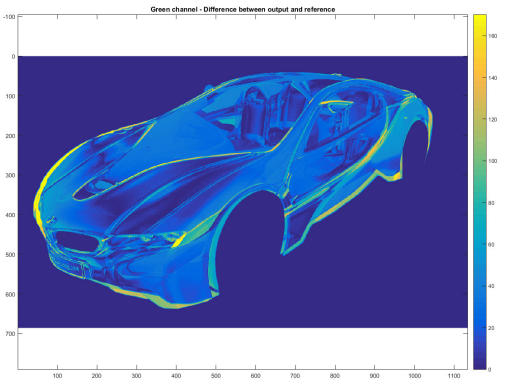


(f) Output - NRMSD Blue Channel 0.1037

Figure 6.19: Comparison of the per-channel error maps before (on the left) and after applying our method (on the right).

# Chapter 7

# Discussion and Conclusions

This project was driven by the need of automotive industry to solve the long standing issue of virtual material consistency: it is a well-known fact that the same material samples tend to differ throughout different renderers, and in different departments different tools are often used, even within the same industry. There are several reason why two different department make use of a different rendering tools, which range from licensing and costs to a difficult integration into their specific workflows, other than the lack of some specific features. Hence, the main goal was to make uniform the appearance of a virtual material across different rendering tools and material models.

## 7.1 Benefits

In order to provide a solution to the lack of consistency in virtual material appearance, the first step was to closely explore BRDF reflectance models and acquisition set ups, in order to gain a deep understanding of the state of the art in Computer Graphics. This work led to the publication "BRDF representation and acquisition" in the prestigious "Computer Graphics Forum journal", presented at Eurographics 2016 as a "State of the Art Report" and two peer-reviewed Courses for the top-venue Conferences "SIGGRAPH Asia 2016" and "SIGGRAPH 2017". This body of research highlighted the inherent difficulties in the virtual material appearance standardisation and allowed us to share our knowledge with other re-

searchers and students in the field, which can now benefit from a clear guidance on the topic of surface representation of material appearance.

The objectives and aims of this thesis, as described in the Introduction Chapter 1, are to mitigate the current tool dependency of material appearance and deliver automatic solutions to the problem of material consistency across renderers, taking into account the constraints of real-world scenarios (such as exporting a material from one software to an external renderer, creating a rendering in a department making use of a specific tool and sending it to a different one, where such tool is not available, having no access to the source code of BRDF models and renderers and so on).

To address reach our goals, possible methods for a general solution to material appearance coherence were hypothesized and tested. The most promising venues have been formalised and defined as minimum and maximum ignorance scenarios (described in Sections 5.1 and 6.1), which are:

(a) information about the parameters used for the source material model and renderer is available (addressed in Chapter 5);

(b) No information about the parameters used for the source material model and renderer (addressed in Chapter 6).

In both cases the aim is the same: make uniform the visual appearance of a rendering obtained with the available material model and renderer (target) to a reference rendering (source) created somewhere else; the main input is the final rendering and renderer used for it.

In particular, when the information about the parameters used for the source material model is available, the proposed automatic image-based framework finds the best set of parameters for the target material model. Given a source BRDF model and a few images of a sphere rendered with it, a genetic algorithm optimisation is able to find a set of parameters in the target BRDF models, such that a sphere rendered using the target model with the optimised parameters matches appearance of the source one. The effectiveness of proposed solution has been evaluated through both numerical validation, user studies and psychometric scaling experiments. Even on scenes with objects with a complex shape, and when

the source and target models have very different characteristics and underlying assumptions, our framework fully meets the research aims stated in the Introduction of this thesis (Section 1.0.4), providing artists with a tool easy to integrate in the typical workflow of digital content creation.

The lack of the information about the parameters used for the source model(s) constitutes a more complex scenario. The proposed framework characterises the typical output of a set of renderers, in order to extract a "visual signature" from the rendered images. If an arbitrary renderer is selected as a reference "golden standard", the framework is able to derive appropriate mapping functions to make the output of a target renderer to look uniform to the reference one, by means of an appearance transfer technique. To be able to describe the effects of a renderer, given an input BRDF model, several classes of materials were included in a dataset test-bench.

A typical scene rendered in the automotive industry includes a wide, but limited, set of materials. Since the developed technique works on a per-material basis, this would leave the task of identifying the material labels within a scene as a manual step, although in some cases the material labels are provided as additional output by renderers. To remove the need of manual material labelling the proposed solution has been augmented by Computer Vision tools, such as Convolutional Neural Network (CNN) and Conditional Random Fields (CRF) from the state of the art, aimed to automatise such task. Overall, the proposed solution constitutes a step forward towards achieving material appearance consistency in real world scenarios. However, there is room for improvement, as discussed in the next session.

## 7.2 Limitations and Future Work

This research provides interesting venues to address in future work. Firstly, our experience with parameters remapping for conductors provided some important insight: the differences between the statistical distribution of the microfacets used by the source and target model seem to play a fundamental role in the quality achievable by parameters remapping techniques, which can be addressed by means of the appearance transfer technique used in the maximum-ignorance sce-

nario. This could lead to an unified and more general framework for consistent material appearance, working at different levels of abstraction (*i.e.* BRDF models parameters and Renderers). Future work includes a more extensive set of psychophysical scaling experiments [177, 178], with a different methodology from the one currently used. In particular, for each set of source BRDF parameters, we plan to perform similarity ranking experiments, in which the observer will be presented with a rendering of the target model with the remapped parameters and a set of variations in which the target parameters are slightly modified; each rendering will display a "blobby" object [179]. Each observer will be requested to rank the rendering from the most similar to the least similar to the source rendering, thus proving either further evidence of the effectiveness of the remapping technique or suggesting a way to correct the fitness function according to the user perception.

The technique for tool independent material appearance can be improved in several ways. The most obvious one would require extending the dataset of reference materials, in order to make it of general application, rather than specifically aimed to the automotive industry. Moreover additional work is required in order to guarantee that the output is physically plausible, thus enforcing all the BRDF properties. In fact, at the current stage, the direct manipulation of BRDF characteristics could lead to some violations of reciprocity and energy conservation.

The size of the additional dataset used to train the Googlenet network [180], described in Section 6.4.3, is far from being big enough to ensure a robust segmentation in all situations. Unfortunately the manual labelling of the images, as described in [171], is extremely time consuming and can only be outsourced to a crowd with tools such as Amazon Mechanical Turk.

The current image-based lighting estimation can be negatively affected by the relative size of the object of interest in the image lattice. An inaccurate lighting estimation can lead to inaccuracies also in the appearance transfer stage, thus leading the hallucinated appearance to deviate from the reference image. A possible way to address this limitation is by means of a dataset of background images, from which the closest one to the actual background in the reference image could be matched, similarly to [157].

Finally, human perception of material appearance could be addressed in order

to identify important visual cues that need to be captured and modelled, in order to include this aspect in the mapping functions and discarding non-perceivable information, thus delivering a perceptually accurate output.

# Bibliography

[1] C. M. Goral, K. E. Torrance, D. P. Greenberg, B. Battaile, Modeling the interaction of light between diffuse surfaces, SIGGRAPH Comput. Graph. 18 (3) (1984) 213–222. doi:10.1145/964965.808601.
URL http://doi.acm.org/10.1145/964965.808601

[2] D. Geisler-Moroder, A. Dr, A new ward brdf model with bounded albedo, Computer Graphics Forum 29 (4) (2010) 1391–1398. doi:10.1111/j.1467-8659.2010.01735.x.
URL http://dx.doi.org/10.1111/j.1467-8659.2010.01735.x

[3] S. Ershov, R. Durikovic, K. Kolchin, K. Myszkowski, Reverse engineering approach to appearance-based design of metallic and pearlescent paints, The Visual Computer 20 (8-9) (2004) 586–600.

[4] W. Matusik, H. Pfister, M. Brand, L. McMillan, A data-driven reflectance model, ACM Trans. Graph. 22 (3) (2003) 759–769. doi:10.1145/882262.882343.
URL http://doi.acm.org/10.1145/882262.882343

[5] D. Guarnera, G. Guarnera, A. Ghosh, C. Denk, M. Glencross, Brdf representation and acquisition, Computer Graphics Forum 35 (2) (2016) 625–650. doi:10.1111/cgf.12867.
URL http://dx.doi.org/10.1111/cgf.12867

[6] E. A. Khan, E. Reinhard, R. W. Fleming, H. H. Bülthoff, Image-based material editing, ACM Trans. Graph. 25 (3) (2006) 654–663.

doi:10.1145/1141911.1141937.
URL http://doi.acm.org/10.1145/1141911.1141937

[7] D. Guarnera, G. C. Guarnera, A. Ghosh, I. Hall, M. Glencross, Capturing and representing brdfs for virtual reality, in: SIGGRAPH ASIA 2016 Courses, SA '16, ACM, New York, NY, USA, 2016, pp. 5:1–5:80. doi:10.1145/2988458.2988508.
URL http://doi.acm.org/10.1145/2988458.2988508

[8] G. C. Guarnera, A. Ghosh, I. Hall, M. Glencross, D. Guarnera, Material capture and representation with applications in virtual reality, in: ACM SIGGRAPH 2017 Courses, SIGGRAPH '17, ACM, New York, NY, USA, 2017, pp. 6:1–6:72. doi:10.1145/3084873.3084918.
URL http://doi.acm.org/10.1145/3084873.3084918

[9] D. Guarnera, G. C. Guarnera, M. Toscani, M. Glencross, B. Li, J. Y. Hardeberg, K. Gegenfurtner, Perceptually validated cross-renderer analytical brdf parameter remapping, *submitted to* IEEE Transactions on Visualization and Computer Graphics.

[10] D. Guarnera, G. C. Guarnera, C. Denk, M. Glencross, Towards a consistent, tool independent virtual material appearance, Electronic Imaging 2017 (8) (2017) 73–79.

[11] I. E. Sutherland, The ultimate display, in: Proceedings of the IFIP Congress, 1965, pp. 506–508.

[12] O. Klehm, F. Rousselle, M. Papas, D. Bradley, C. Hery, B. Bickel, W. Jarosz, T. Beeler, Recent advances in facial appearance capture, in: Computer Graphics Forum, Vol. 34, Wiley Online Library, 2015, pp. 709–733.

[13] B. L. Anderson, Visual perception of materials and surfaces, Current Biology 21 (24) (2011) R978 – R983. doi:https://doi.org/10.1016/j.cub.2011.11.022.

[14] R. W. Fleming, Visual perception of materials and their properties, Vision Research 94 (Supplement C) (2014) 62 – 75. doi:https://doi.org/10.1016/j.visres.2013.11.004.

[15] B. L. Anderson, J. Kim, Image statistics do not explain the perception of gloss and lightness, Journal of vision 9 (11) (2009) 10–10.

[16] J. Kim, B. L. Anderson, Image statistics and the perception of surface gloss and lightness, Journal of Vision 10 (9) (2010) 3–3.

[17] C. Witzel, J. K. O'Regan, S. Hansmann-Roth, The dress and individual differences in the perception of surface properties, Vision Research-doi:https://doi.org/10.1016/j.visres.2017.07.015.

[18] J. Krauskopf, G. Karl, Color discrimination and adaptation, Vision research 32 (11) (1992) 2165–2175.

[19] H. W. Jensen, S. R. Marschner, M. Levoy, P. Hanrahan, A practical model for subsurface light transport, in: Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH '01, ACM, New York, NY, USA, 2001, pp. 511–518. doi:10.1145/383259.383319.
URL `http://doi.acm.org/10.1145/383259.383319`

[20] F. Nicodemus, J. Richmond, J. Hsia, I. Ginsberg, T. Limperis, Geometrical considerations and nomenclature for reflectance, natl, Bur. Stand. Rep., NBS MN-160.

[21] M. Goesele, H. P. A. Lensch, J. Lang, C. Fuchs, H.-P. Seidel, Disco: Acquisition of translucent objects, ACM Trans. Graph. 23 (3) (2004) 835–844. doi:10.1145/1015706.1015807.
URL `http://doi.acm.org/10.1145/1015706.1015807`

[22] C. Dachsbacher, M. Stamminger, Translucent shadow maps, in: Proceedings of the 14th Eurographics Workshop on Rendering, EGRW '03, Eurographics Association, Aire-la-Ville, Switzerland, Switzerland, 2003, pp.

197–201.
URL http://dl.acm.org/citation.cfm?id=882404.882433

[23] X. Hao, T. Baby, A. Varshney, Interactive subsurface scattering for translucent meshes, in: Proceedings of the 2003 Symposium on Interactive 3D Graphics, I3D '03, ACM, New York, NY, USA, 2003, pp. 75–82. doi:10.1145/641480.641497.
URL http://doi.acm.org/10.1145/641480.641497

[24] C. Donner, T. Weyrich, E. d'Eon, R. Ramamoorthi, S. Rusinkiewicz, A layered, heterogeneous reflectance model for acquiring and rendering human skin, ACM Trans. Graph. 27 (5) (2008) 140:1–140:12. doi:10.1145/1409060.1409093.
URL http://doi.acm.org/10.1145/1409060.1409093

[25] E. D'Eon, G. Irving, A quantized-diffusion model for rendering translucent materials, ACM Trans. Graph. 30 (4) (2011) 56:1–56:14. doi:10.1145/2010324.1964951.
URL http://doi.acm.org/10.1145/2010324.1964951

[26] O. Klehm, F. Rousselle, M. Papas, D. Bradley, C. Hery, B. Bickel, W. Jarosz, T. Beeler, Recent advances in facial appearance capture, Computer Graphics Forum (Proceedings of Eurographics) 34 (2) (2015) 709–733. doi:10.1111/cgf.12594.

[27] S. Rusinkiewicz, A new change of variables for efficient brdf representation, in: G. Drettakis, N. Max (Eds.), Rendering Techniques 98, Eurographics Book Series, Springer-Verlag, Wien, Austria, 1998, proceedings of the Workshop held in Vienna, Austria, between Jyne 29th and July 1st 1998.

[28] M. Stark, J. Arvo, B. Smits, Barycentric parameterizations for isotropic brdfs, Visualization and Computer Graphics, IEEE Transactions on 11 (2) (2005) 126–138. doi:10.1109/TVCG.2005.26.

[29] P. Barla, L. Belcour, R. Pacanowski, In praise of an alternative brdf parametrization, in: Workshop on Material Appearance Modeling , Pro-

ceedings of the Workshop on Material Appearance Modeling 2015, Darmstadt, Germany, 2015.
URL https://hal.inria.fr/hal-01172118

[30] F. Clarke, D. Parry, Helmholtz reciprocity: its validity and application to reflectometry, Lighting Research and Technology 17 (1) (1985) 1–11. doi:10.1177/14771535850170010301.

[31] E. Veach, Robust monte carlo methods for light transport simulation., Ph.D. thesis, Stanford University (1997).

[32] J. Dorsey, H. Rushmeier, F. Sillion, Digital Modeling of Material Appearance, Morgan Kaufmann, 2007.

[33] C. Schlick, An inexpensive brdf model for physically-based rendering, Computer Graphics Forum 13 (3) (1994) 233–246. doi:10.1111/1467-8659.1330233.
URL http://dx.doi.org/10.1111/1467-8659.1330233

[34] J. T. Kajiya, T. L. Kay, Rendering fur with three dimensional textures, SIGGRAPH Comput. Graph. 23 (3) (1989) 271–280. doi:10.1145/74334.74361.
URL http://doi.acm.org/10.1145/74334.74361

[35] K. Perlin, E. M. Hoffert, Hypertexture, SIGGRAPH Comput. Graph. 23 (3) (1989) 253–262. doi:10.1145/74334.74359.
URL http://doi.acm.org/10.1145/74334.74359

[36] Y.-Q. Xu, Y. Chen, S. Lin, H. Zhong, E. Wu, B. Guo, H.-Y. Shum, Photorealistic rendering of knitwear using the lumislice, in: Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH '01, ACM, New York, NY, USA, 2001, pp. 391–398. doi:10.1145/383259.383303.
URL http://doi.acm.org/10.1145/383259.383303

[37] W. Jakob, A. Arbree, J. T. Moon, K. Bala, S. Marschner, A radiative transfer framework for rendering materials with anisotropic structure, ACM Trans. Graph. 29 (4) (2010) 53:1–53:13.

[38] S. Zhao, W. Jakob, S. Marschner, K. Bala, Building volumetric appearance models of fabric using micro ct imaging, ACM Trans. Graph. 30 (4) (2011) 44:1–44:10. doi:10.1145/2010324.1964939.
URL `http://doi.acm.org/10.1145/2010324.1964939`

[39] P. Khungurn, D. Schroeder, S. Zhao, K. Bala, S. Marschner, Matching real fabrics with micro-appearance models, ACM Trans. Graph. 35 (1) (2015) 1:1–1:26. doi:10.1145/2818648.
URL `http://doi.acm.org/10.1145/2818648`

[40] E. A. Haines, Beams o' light: Confessions of a hacker, SIGGRAPH 91 Course Notes - Frontiers in Rendering.

[41] J. Lawrence, S. Rusinkiewicz, R. Ramamoorthi, Efficient brdf importance sampling using a factored representation, in: ACM Transactions on Graphics (TOG), Vol. 23, ACM, 2004, pp. 496–505.

[42] P. Clarberg, W. Jarosz, T. Akenine-Möller, H. W. Jensen, Wavelet importance sampling: efficiently evaluating products of complex functions, in: ACM Transactions on Graphics (TOG), Vol. 24, ACM, 2005, pp. 1166–1175.

[43] B. T. Phong, Illumination for computer generated pictures, Commun. ACM 18 (6) (1975) 311–317. doi:10.1145/360825.360839.
URL `http://doi.acm.org/10.1145/360825.360839`

[44] W. Jakob, Mitsuba renderer, http://www.mitsuba-renderer.org (2010).

[45] J. Arvo, Applications of irradiance tensors to the simulation of non-lambertian phenomena, in: Proceedings of the 22Nd Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH '95, ACM, New York, NY, USA, 1995, pp. 335–342. doi:10.1145/218380.218467.
URL `http://doi.acm.org/10.1145/218380.218467`

[46] R. R. Lewis, Making shaders more physically plausible, Computer Graphics Forum 13 (2) (1994) 109–120. doi:10.1111/1467-8659.1320109.
URL http://dx.doi.org/10.1111/1467-8659.1320109

[47] J. F. Blinn, Models of light reflection for computer synthesized pictures, SIGGRAPH Comput. Graph. 11 (2) (1977) 192–198. doi:10.1145/965141.563893.
URL http://doi.acm.org/10.1145/965141.563893

[48] A. Brady, J. Lawrence, P. Peers, W. Weimer, genbrdf: Discovering new analytic brdfs with genetic programming, ACM Trans. Graph. 33 (4) (2014) 114:1–114:11. doi:10.1145/2601097.2601193.

[49] G. J. Ward, Measuring and modeling anisotropic reflection, SIGGRAPH Comput. Graph. 26 (2) (1992) 265–272. doi:10.1145/142920.134078.
URL http://doi.acm.org/10.1145/142920.134078

[50] L. Neumann, A. Neumannn, L. Szirmay-Kalos, Compact metallic reflectance models, Computer Graphics Forum 18 (3) (1999) 161–172. doi:10.1111/1467-8659.00337.
URL http://dx.doi.org/10.1111/1467-8659.00337

[51] A. Dür, An improved normalization for the ward reflectance model, Journal of Graphics, GPU, and Game Tools 11 (1) (2006) 51–59. doi:10.1080/2151237X.2006.10129215.

[52] E. P. F. Lafortune, S.-C. Foo, K. E. Torrance, D. P. Greenberg, Non-linear approximation of reflectance functions, in: Proceedings of the 24th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH '97, ACM Press/Addison-Wesley Publishing Co., New York, NY, USA, 1997, pp. 117–126. doi:10.1145/258734.258801.
URL http://dx.doi.org/10.1145/258734.258801

[53] S. H. Westin, H. Li, K. E. Torrance, A comparison of four brdf models, in: Eurographics Symposium on Rendering. pags, 2004, pp. 1–10.

[54] E. P. Lafortune, Y. D. Willems, Using the modified phong reflectance model for physically based rendering, Tech. rep. (1994).

[55] M. Ashikhmin, P. Shirley, An anisotropic phong brdf model, J. Graph. Tools 5 (2) (2000) 25–32. doi:10.1080/10867651.2000.10487522.
URL http://dx.doi.org/10.1080/10867651.2000.10487522

[56] A. Ozturk, M. Kurt, A. Bilgili, C. Gungor, Linear approximation of bidirectional reflectance distribution functions, Computers & Graphics 32 (2) (2008) 149–158.

[57] D. Edwards, S. Boulos, J. Johnson, P. Shirley, M. Ashikhmin, M. Stark, C. Wyman, The halfway vector disk for brdf modeling, ACM Trans. Graph. 25 (1) (2006) 1–18. doi:10.1145/1122501.1122502.
URL http://doi.acm.org/10.1145/1122501.1122502

[58] M. Ashikhmin, S. Premoze, Distribution-based brdfs, Tech. rep. (2007).

[59] K. E. Torrance, E. M. Sparrow, Theory for off-specular reflection from roughened surfaces, JOSA 57 (9) (1967) 1105–1112.

[60] R. L. Cook, K. E. Torrance, A reflectance model for computer graphics, ACM Transactions on Graphics (TOG) 1 (1) (1982) 7–24.

[61] A. Ghosh, T. Hawkins, P. Peers, S. Frederiksen, P. Debevec, Practical modeling and acquisition of layered facial reflectance, in: ACM Transactions on Graphics (TOG), Vol. 27, ACM, 2008, p. 139.

[62] T. Akenine-Möller, E. Haines, N. Hoffman, Real-Time Rendering 3rd Edition, A. K. Peters, Ltd., Natick, MA, USA, 2008.

[63] E. Heitz, Understanding the masking-shadowing function in microfacet-based brdfs, Journal of Computer Graphics Techniques (JCGT) 3 (2) (2014) 48–107.
URL http://jcgt.org/published/0003/02/03/

[64] T. Weyrich, J. Lawrence, H. Lensch, S. Rusinkiewicz, T. Zickler, Principles of appearance acquisition and representation, Foundations and Trends® in Computer Graphics and Vision 4 (2) (2009) 75–191.

[65] X. D. He, K. E. Torrance, F. X. Sillion, D. P. Greenberg, A comprehensive physical model for light reflection, SIGGRAPH Comput. Graph. 25 (4) (1991) 175–186. doi:10.1145/127719.122738.
URL `http://doi.acm.org/10.1145/127719.122738`

[66] A. Ngan, F. Durand, W. Matusik, Experimental analysis of brdf models, in: Proceedings of the Sixteenth Eurographics Conference on Rendering Techniques, EGSR'05, Eurographics Association, Aire-la-Ville, Switzerland, Switzerland, 2005, pp. 117–126. doi:10.2312/EGWR/EGSR05/117-126.

[67] M. Oren, S. K. Nayar, Generalization of lambert's reflectance model, in: Proceedings of the 21st annual conference on Computer graphics and interactive techniques, ACM, 1994, pp. 239–246.

[68] S. Ershov, K. Kolchin, K. Myszkowski, Rendering pearlescent appearance based on paint-composition modelling, in: Computer Graphics Forum, Vol. 20, Wiley Online Library, 2001, pp. 227–238.

[69] B. Walter, S. R. Marschner, H. Li, K. E. Torrance, Microfacet models for refraction through rough surfaces, in: Proceedings of the 18th Eurographics conference on Rendering Techniques, Eurographics Association, 2007, pp. 195–206.

[70] B. Smith, Geometrical shadowing of a random rough surface, Antennas and Propagation, IEEE Transactions on 15 (5) (1967) 668–671.

[71] M. M. Bagher, C. Soler, N. Holzschuch, Accurate fitting of measured reflectances using a shifted gamma micro-facet distribution, in: Computer Graphics Forum, Vol. 31, Wiley Online Library, 2012, pp. 1509–1518.

[72] T. S. Trowbridge, K. P. Reitz, Average irregularity representation of a rough surface for ray reflection, J. Opt. Soc. Am. 65 (5) (1975) 531–536. doi:10.1364/JOSA.65.000531.

[73] S. McAuley, S. Hill, N. Hoffman, Y. Gotanda, B. Smits, B. Burley, A. Martinez, Practical physically-based shading in film and game production, in: ACM SIGGRAPH 2012 Courses, SIGGRAPH '12, ACM, New York, NY, USA, 2012, pp. 10:1–10:7. doi:10.1145/2343483.2343493.
URL http://doi.acm.org/10.1145/2343483.2343493

[74] E. Heitz, J. Dupuy, C. Crassin, C. Dachsbacher, The sggx microflake distribution, ACM Trans. Graph. 34 (4) (2015) 48:1–48:11. doi:10.1145/2766988.
URL http://doi.acm.org/10.1145/2766988

[75] M. Rump, G. Müller, R. Sarlette, D. Koch, R. Klein, Photo-realistic rendering of metallic car paint from image-based measurements, Computer Graphics Forum 27 (2) (2008) 527–536.

[76] M. Kurt, L. Szirmay-Kalos, J. Křivánek, An anisotropic brdf model for fitting and monte carlo rendering, SIGGRAPH Comput. Graph. 44 (1) (2010) 3:1–3:15. doi:10.1145/1722991.1722996.
URL http://doi.acm.org/10.1145/1722991.1722996

[77] J. Löw, J. Kronander, A. Ynnerman, J. Unger, Brdf models for accurate and efficient rendering of glossy surfaces, ACM Transactions on Graphics (TOG) 31 (1) (2012) 9.

[78] E. L. Church, P. Z. Takacs, T. A. Leonard, The prediction of brdfs from surface profile measurements, in: Proc. SPIE, Vol. 1165, 1990, pp. 136–150. doi:10.1117/12.962842.
URL http://dx.doi.org/10.1117/12.962842

[79] E. L. Church, P. Z. Takacs, Optimal estimation of finish parameters, in: Proc. SPIE, Vol. 1530, 1991, pp. 71–85. doi:10.1117/12.50498.
URL http://dx.doi.org/10.1117/12.50498

[80] W. Jakob, M. Hašan, L.-Q. Yan, J. Lawrence, R. Ramamoorthi, S. Marschner, Discrete stochastic microfacet models, ACM Trans. Graph. 33 (4) (2014) 115:1–115:10. doi:10.1145/2601097.2601186.
URL http://doi.acm.org/10.1145/2601097.2601186

[81] A. Weidlich, A. Wilkie, Arbitrarily layered micro-facet surfaces, in: Proceedings of the 5th International Conference on Computer Graphics and Interactive Techniques in Australia and Southeast Asia, GRAPHITE '07, ACM, New York, NY, USA, 2007, pp. 171–178. doi:10.1145/1321261.1321292.
URL http://doi.acm.org/10.1145/1321261.1321292

[82] J. Dupuy, E. Heitz, J.-C. Iehl, P. Poulin, V. Ostromoukhov, Extracting microfacet-based brdf parameters from arbitrary materials with power iterations, Computer Graphics Forum 34 (4) (2015) 21–30. doi:10.1111/cgf.12675.
URL http://dx.doi.org/10.1111/cgf.12675

[83] J. Filip, M. Havlíček, R. Vávra, Adaptive highlights stencils for modeling of multi-axial brdf anisotropy, The Visual Computer (2015) 1–11doi:10.1007/s00371-015-1148-1.
URL http://dx.doi.org/10.1007/s00371-015-1148-1

[84] P. Lalonde, A. Fournier, A wavelet representation of reflectance functions, IEEE Transactions on Visualization and Computer Graphics 3 (4) (1997) 329–336. doi:10.1109/2945.646236.
URL http://dx.doi.org/10.1109/2945.646236

[85] L. Claustres, M. Paulin, Y. Boucher, Brdf measurement modelling using wavelets for efficient path tracing, Computer Graphics Forum (2003) 701–716.

[86] L. Claustres, Y. Boucher, M. Paulin, Wavelet-based modeling of spectral bidirectional reflectance distribution function data, Optical Engineering 43 (10) (2004) 2327–2339.

[87] S. H. Westin, J. R. Arvo, K. E. Torrance, Predicting reflectance functions from complex surfaces, SIGGRAPH Comput. Graph. 26 (2) (1992) 255–264. doi:10.1145/142920.134075.
URL `http://doi.acm.org/10.1145/142920.134075`

[88] W. H. Press, B. P. Flannery, S. A. Teukolsky, W. T. Vetterling, Numerical Recipes in C: The Art of Scientific Computing, Cambridge University Press, New York, NY, USA, 1988.

[89] A. Fournier, Separating reflection functions for linear radiosity, in: Rendering Techniques 95, Eurographics, Springer Vienna, 1995, pp. 296–305.

[90] M. Brand, Charting a manifold, in: Advances in Neural Information Processing Systems 15, MIT Press, 2002, pp. 961–968.

[91] W. Matusik, H. Pfister, M. Brand, L. McMillan, Efficient isotropic brdf measurement, in: Proceedings of the 14th Eurographics Workshop on Rendering, EGRW '03, Eurographics Association, Aire-la-Ville, Switzerland, Switzerland, 2003, pp. 241–247.
URL `http://dl.acm.org/citation.cfm?id=882404.882439`

[92] F. Romeiro, Y. Vasilyev, T. Zickler, Passive reflectometry, in: Proceedings of the 10th European Conference on Computer Vision: Part IV, ECCV '08, Springer-Verlag, Berlin, Heidelberg, 2008, pp. 859–872.

[93] J. Kautz, M. D. McCool, Interactive rendering with arbitrary brdfs using separable approximations, in: Rendering Techniques 99, Springer, 1999, pp. 247–260.

[94] M. D. McCool, J. Ang, A. Ahmad, Homomorphic factorization of brdfs for high-performance rendering, in: Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH '01, ACM, New York, NY, USA, 2001, pp. 171–178. doi:10.1145/383259.383276.

[95] J. Lawrence, A. Ben-Artzi, C. DeCoro, W. Matusik, H. Pfister, R. Ramamoorthi, S. Rusinkiewicz, Inverse shade trees for non-parametric ma-

terial representation and editing, in: ACM SIGGRAPH 2006 Papers, SIGGRAPH '06, ACM, New York, NY, USA, 2006, pp. 735–745. doi:10.1145/1179352.1141949.

[96] X. Sun, K. Zhou, Y. Chen, S. Lin, J. Shi, B. Guo, Interactive relighting with dynamic brdfs, ACM Trans. Graph. 26 (3). doi:10.1145/1276377.1276411.
URL http://doi.acm.org/10.1145/1276377.1276411

[97] A. Bilgili, A. Öztürk, M. Kurt, A general brdf representation based on tensor decomposition, Computer Graphics Forum 30 (8) (2011) 2427–2439. doi:10.1111/j.1467-8659.2011.02072.x.
URL http://dx.doi.org/10.1111/j.1467-8659.2011.02072.x

[98] L. Tucker, Some mathematical notes on three-mode factor analysis, Psychometrika 31 (3) (1966) 279–311. doi:10.1007/BF02289464.
URL http://dx.doi.org/10.1007/BF02289464

[99] R. Pacanowski, O. Salazar Celis, C. Schlick, X. Granier, P. Poulin, A. Cuyt, Rational brdf, Visualization and Computer Graphics, IEEE Transactions on 18 (11) (2012) 1824–1835. doi:10.1109/TVCG.2012.73.

[100] G. Ward, M. Kurt, N. Bonneel, Framework for sharing and rendering real-world bidirectional scattering distribution functions, Tech. Rep. LBNL-5954E, Lawrence Berkeley National Laboratory (October 2012).

[101] G. Ward, M. Kurt, N. Bonneel, Reducing anisotropic bsdf measurement to common practice, in: Proceedings of the Eurographics 2014 Workshop on Material Appearance Modeling: Issues and Acquisition, MAM '14, Eurographics Association, Aire-la-Ville, Switzerland, Switzerland, 2014, pp. 5–8.

[102] K. Nishino, S. Lombardi, Directional statistics-based reflectance model for isotropic bidirectional reflectance distribution functions, J. Opt. Soc. Am. A 28 (1) (2011) 8–18. doi:10.1364/JOSAA.28.000008.
URL http://josaa.osa.org/abstract.cfm?URI=josaa-28-1-8

[103] M. Ashikhmin, P. Shirley, S. Marschner, J. Stam, State of the art in modeling and measuring of surface reflection, in: ACM SIGGRAPH 2001 Courses, ACM, 2001, p. 1.

[104] J. Y. Han, K. Perlin, Measuring bidirectional texture reflectance with a kaleidoscope, ACM Trans. Graph. 22 (3) (2003) 741–748. doi:10.1145/882262.882341.
URL http://doi.acm.org/10.1145/882262.882341

[105] H. Rushmeier, Y. D. Lockerman, L. Cartwright, D. Pitera, Experiments with a low-cost system for computer graphics material model acquisition, Measuring, Modeling, and Reproducing Material Appearance 2015, Measuring, Modeling, and Reproducing Material Appearance 2015, 2015.

[106] R. J. Woodham, Photometric method for determining surface orientation from multiple images, Optical Engineering 19 (1) (1980) 191139–191139–. doi:10.1117/12.7972479.
URL http://dx.doi.org/10.1117/12.7972479

[107] A. S. Georghiades, Recovering 3-d shape and reflectance from a small number of photographs, in: Proceedings of the 14th Eurographics Workshop on Rendering, EGRW '03, Eurographics Association, Aire-la-Ville, Switzerland, Switzerland, 2003, pp. 230–240.
URL http://dl.acm.org/citation.cfm?id=882404.882438

[108] D. B. Goldman, B. Curless, A. Hertzmann, S. M. Seitz, Shape and spatially-varying brdfs from photometric stereo, in: Proceedings of the Tenth IEEE International Conference on Computer Vision (ICCV'05) Volume 1 - Volume 01, ICCV '05, IEEE Computer Society, Washington, DC, USA, 2005, pp. 341–348. doi:10.1109/ICCV.2005.219.
URL http://dx.doi.org/10.1109/ICCV.2005.219

[109] A. Hertzmann, S. M. Seitz, Example-based photometric stereo: Shape reconstruction with general, varying brdfs, Pattern Analysis and Machine Intelligence, IEEE Transactions on 27 (8) (2005) 1254–1264.

[110] T. Chen, M. Goesele, H.-P. Seidel, Mesostructure from specularity, in: Computer Vision and Pattern Recognition, 2006 IEEE Computer Society Conference on, Vol. 2, IEEE, 2006, pp. 1825–1832.

[111] T. Zickler, R. Ramamoorthi, S. Enrique, P. N. Belhumeur, Reflectance sharing: Predicting appearance from a sparse set of images of a known shape, IEEE Trans. Pattern Anal. Mach. Intell. 28 (8) (2006) 1287–1302. doi:10.1109/TPAMI.2006.170.
URL `http://dx.doi.org/10.1109/TPAMI.2006.170`

[112] N. Alldrin, T. Zickler, D. Kriegman, Photometric stereo with non-parametric and spatially-varying reflectance, in: Computer Vision and Pattern Recognition, 2008. CVPR 2008. IEEE Conference on, IEEE, 2008, pp. 1–8.

[113] M. Holroyd, J. Lawrence, G. Humphreys, T. Zickler, A photometric approach for estimating normals and tangents, ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia 2008) 27 (5).

[114] M. Goesele, H. Lensch, H.-P. Seidel, Validation of color managed 3d appearance acquisition, in: Color and Imaging Conference, Vol. 2004, Society for Imaging Science and Technology, 2004, pp. 265–270.

[115] G. Wyszecki, W. S. Stiles, Color science: Concepts and Methods, Quantitative Data and Formulae, Vol. 8, Wiley, New York, NY, USA, 1982.

[116] M. J. Fairchild, Color Appearance Models, John Wiley & Sons, 2005.

[117] G. Ward, E. Eydelberg-Vileshin, Picture perfect rgb rendering using spectral prefiltering and sharp color primaries, in: Proceedings of the 13th Eurographics Workshop on Rendering, EGRW '02, Eurographics Association, Aire-la-Ville, Switzerland, Switzerland, 2002, pp. 117–124.
URL `http://dl.acm.org/citation.cfm?id=581896.581913`

[118] M. H. Kim, J. Kautz, Characterization for high dynamic range imaging, Computer Graphics Forum 27 (2) (2008) 691–697. doi:10.1111/j.1467-

8659.2008.01167.x.
URL http://dx.doi.org/10.1111/j.1467-8659.2008.01167.x

[119] G. C. Guarnera, S. Bianco, R. Schettini, Absolute colorimetric characterization of a dslr camera, in: Proc. SPIE, Vol. 9023, 2014, pp. 90230U–90230U–7. doi:10.1117/12.2042172.
URL http://dx.doi.org/10.1117/12.2042172

[120] B. Cabral, N. Max, R. Springmeyer, Bidirectional reflection functions from surface bump maps, SIGGRAPH Comput. Graph. 21 (4) (1987) 273–281. doi:10.1145/37402.37434.
URL http://doi.acm.org/10.1145/37402.37434

[121] P. Hanrahan, W. Krueger, Reflection from layered surfaces due to subsurface scattering, in: Proceedings of the 20th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH '93, ACM, New York, NY, USA, 1993, pp. 165–174. doi:10.1145/166117.166139.
URL http://doi.acm.org/10.1145/166117.166139

[122] M. Ashikmin, S. Premože, P. Shirley, A microfacet-based brdf generator, in: Proceedings of the 27th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH '00, ACM Press/Addison-Wesley Publishing Co., New York, NY, USA, 2000, pp. 65–74. doi:10.1145/344779.344814.
URL http://dx.doi.org/10.1145/344779.344814

[123] Z. Dong, B. Walter, S. Marschner, D. P. Greenberg, Predicting appearance from measured microgeometry of metal surfaces, ACM Trans. Graph. 35 (1) (2015) 9:1–9:13. doi:10.1145/2815618.
URL http://doi.acm.org/10.1145/2815618

[124] J. J. Hsia, J. C. Richmond, Bidirectional reflectometry. part i: A high resolution laser bidirectional reflectometer with results on several optical coatings, Journal of Reseach of the National Bureau of Standards-A. Physics and Chemistry A 80 (1976) 189–205.

[125] S. C. Foo, A gonioreflectometer for measuring the bidirectional reflectance of material for use in illumination computation, Ph.D. thesis, Cornell University (1997).

[126] H. Li, S.-C. Foo, K. E. Torrance, S. H. Westin, Automated three-axis gonioreflectometer for computer graphics applications, Optical Engineering 45 (4) (2006) 043605–043605.

[127] N. Riviere, R. Ceolato, L. Hespel, Multispectral polarized brdf: design of a highly resolved reflectometer and development of a data inversion method, Opt. Appl 42.

[128] S. R. Marschner, P. D, Inverse rendering for computer graphics, Tech. rep. (1998).

[129] S. R. Marschner, S. H. Westin, E. P. Lafortune, K. E. Torrance, D. P. Greenberg, Image-based brdf measurement including human skin, in: Rendering Techniques 99, Springer, 1999, pp. 131–144.

[130] M. Kurt, D. Edwards, A survey of brdf models for computer graphics, ACM SIGGRAPH Computer Graphics 43 (2) (2009) 4.

[131] N. Naik, S. Zhao, A. Velten, R. Raskar, K. Bala, Single view reflectance capture using multiplexed scattering and time-of-flight imaging, ACM Trans. Graph. 30 (6) (2011) 171:1–171:10. doi:10.1145/2070781.2024205. URL `http://doi.acm.org/10.1145/2070781.2024205`

[132] K. J. Dana, B. Van Ginneken, S. K. Nayar, J. J. Koenderink, Reflectance and texture of real-world surfaces, ACM Transactions on Graphics (TOG) 18 (1) (1999) 1–34.

[133] Y. Mukaigawa, K. Sumino, Y. Yagi, Multiplexed illumination for measuring brdf using an ellipsoidal mirror and a projector, in: Computer Vision–ACCV 2007, Springer, 2007, pp. 246–257.

[134] A. Ghosh, S. Achutha, W. Heidrich, M. O'Toole, Brdf acquisition with basis illumination, in: Computer Vision, 2007. ICCV 2007. IEEE 11th International Conference on, 2007, pp. 1–8. doi:10.1109/ICCV.2007.4408935.

[135] A. Ghosh, W. Heidrich, S. Achutha, M. O′Toole, A basis illumination approach to brdf measurement, International journal of computer vision 90 (2) (2010) 183–197.

[136] T. Malzbender, D. Gelb, H. Wolters, Polynomial texture maps, in: Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH '01, ACM, New York, NY, USA, 2001, pp. 519–528. doi:10.1145/383259.383320.

[137] M. Ben-Ezra, J. Wang, B. Wilburn, X. Li, L. Ma, An led-only brdf measurement device, in: Computer Vision and Pattern Recognition, 2008. CVPR 2008. IEEE Conference on, 2008, pp. 1–8. doi:10.1109/CVPR.2008.4587766.

[138] A. Ghosh, T. Chen, P. Peers, C. A. Wilson, P. Debevec, Estimating specular roughness and anisotropy from second order spherical gradient illumination, in: Proceedings of the Twentieth Eurographics Conference on Rendering, EGSR'09, Eurographics Association, Aire-la-Ville, Switzerland, Switzerland, 2009, pp. 1161–1170. doi:10.1111/j.1467-8659.2009.01493.x.

[139] W.-C. Ma, T. Hawkins, P. Peers, C.-F. Chabert, M. Weiss, P. Debevec, Rapid acquisition of specular and diffuse normal maps from polarized spherical gradient illumination, in: Proceedings of the 18th Eurographics Conference on Rendering Techniques, EGSR'07, Eurographics Association, Aire-la-Ville, Switzerland, Switzerland, 2007, pp. 183–194. doi:10.2312/EGWR/EGSR07/183-194.
URL `http://dx.doi.org/10.2312/EGWR/EGSR07/183-194`

[140] A. Ghosh, T. Chen, P. Peers, C. A. Wilson, P. Debevec, Circularly polarized spherical illumination reflectometry, ACM Trans. Graph. 29 (6) (2010) 162:1–162:12. doi:10.1145/1882261.1866163.

[141] G. Guarnera, P. Peers, P. Debevec, A. Ghosh, Estimating surface normals from spherical stokes reflectance fields, in: A. Fusiello, V. Murino, R. Cucchiara (Eds.), Computer Vision - ECCV 2012. Workshops and Demonstra-

tions, Vol. 7584 of Lecture Notes in Computer Science, Springer Berlin Heidelberg, 2012, pp. 340–349.

[142] B. Tunwattanapong, G. Fyffe, P. Graham, J. Busch, X. Yu, A. Ghosh, P. Debevec, Acquiring reflectance and shape from continuous spherical harmonic illumination, ACM Trans. Graph. 32 (4) (2013) 109:1–109:12.
URL http://doi.acm.org/10.1145/2461912.2461944

[143] A. Gardner, C. Tchou, T. Hawkins, P. Debevec, Linear light source reflectometry, ACM Trans. Graph. 22 (3) (2003) 749–758.
URL http://doi.acm.org/10.1145/882262.882342

[144] P. Ren, J. Wang, J. Snyder, X. Tong, B. Guo, Pocket reflectometry, ACM Trans. Graph. 30 (4) (2011) 45:1–45:10. doi:10.1145/2010324.1964940.

[145] G. Chen, Y. Dong, P. Peers, J. Zhang, X. Tong, Reflectance scanning: Estimating shading frame and brdf with generalized linear light sources, ACM Trans. Graph. 33 (4) (2014) 117:1–117:11.
URL http://doi.acm.org/10.1145/2601097.2601180

[146] Y. Francken, T. Cuypers, T. Mertens, J. Gielis, P. Bekaert, High quality mesostructure acquisition using specularities, in: CVPR, 2008.

[147] M. Aittala, T. Weyrich, J. Lehtinen, Practical svbrdf capture in the frequency domain, ACM Trans. Graph. 32 (4) (2013) 110:1–110:12.
URL http://doi.acm.org/10.1145/2461912.2461978

[148] C.-P. Wang, N. Snavely, S. Marschner, Estimating dual-scale properties of glossy surfaces from step-edge lighting, in: Proceedings of the 2011 SIGGRAPH Asia Conference, SA '11, ACM, New York, NY, USA, 2011, pp. 172:1–172:12.
URL http://doi.acm.org/10.1145/2024156.2024206

[149] J. Riviere, P. Peers, A. Ghosh, Mobile surface reflectometry, Computer Graphics Forum (2015) n/a–n/adoi:10.1111/cgf.12719.
URL http://dx.doi.org/10.1111/cgf.12719

[150] M. Aittala, T. Weyrich, J. Lehtinen, Two-shot svbrdf capture for stationary materials, ACM Trans. Graph. 34 (4) (2015) 110:1–110:13.
URL http://doi.acm.org/10.1145/2766967

[151] H. P. A. Lensch, J. Kautz, M. Goesele, W. Heidrich, H.-P. Seidel, Image-based reconstruction of spatial appearance and geometric detail, ACM Trans. Graph. 22 (2) (2003) 234–257. doi:10.1145/636886.636891.
URL http://doi.acm.org/10.1145/636886.636891

[152] Y. Dong, J. Wang, X. Tong, J. Snyder, Y. Lan, M. Ben-Ezra, B. Guo, Manifold bootstrapping for svbrdf capture, ACM Trans. Graph. 29 (4) (2010) 98:1–98:10.
URL http://doi.acm.org/10.1145/1778765.1778835

[153] C. J. Zubiaga, A. Muñoz, L. Belcour, C. Bosch, P. Barla, MatCap Decomposition for Dynamic Appearance Manipulation, in: Eurographics Symposium on Rendering 2015, Darmstadt, Germany, 2015.

[154] C. J. Zubiaga, L. Belcour, C. Bosch, A. Muñoz, P. Barla, Statistical analysis of bidirectional reflectance distribution functions, in: SPIE/IS&T Electronic Imaging, International Society for Optics and Photonics, 2015, pp. 939808–939808.

[155] A. Serrano, D. Gutierrez, K. Myszkowski, H.-P. Seidel, B. Masia, An intuitive control space for material appearance, ACM Trans. Graph. 35 (6) (2016) 186:1–186:12. doi:10.1145/2980179.2980242.
URL http://doi.acm.org/10.1145/2980179.2980242

[156] A. Sztrajman, J. Křivánek, A. Wilkie, T. Weyrich, Image-based remapping of material appearance, in: Proceedings of Workshop on Material Appearance Modeling, 2017.

[157] Y. Shih, S. Paris, F. Durand, W. T. Freeman, Data-driven hallucination of different times of day from a single outdoor photo, ACM Trans. Graph. 32 (6) (2013) 200:1–200:11.

[158] T. Günther, T. Grosch, Consistent scene editing by progressive difference images, Comput. Graph. Forum 34 (4) (2015) 41–51. doi:10.1111/cgf.12677.
URL http://dx.doi.org/10.1111/cgf.12677

[159] R. Rojas, Neural Networks - A Systematic Introduction, Springer Berlin Heidelberg, 1996.

[160] R. K. Mantiuk, A. Tomaszewska, R. Mantiuk, Comparison of four subjective methods for image quality assessment, Comput. Graph. Forum 31 (8) (2012) 2478–2491. doi:10.1111/j.1467-8659.2012.03188.x.
URL http://dx.doi.org/10.1111/j.1467-8659.2012.03188.x

[161] M. Glencross, G. J. Ward, F. Melendez, C. Jay, J. Liu, R. Hubbold, A perceptually validated model for surface depth hallucination, ACM Trans. Graph. 27 (3) (2008) 59:1–59:8. doi:10.1145/1360612.1360658.
URL http://doi.acm.org/10.1145/1360612.1360658

[162] G. Ward, M. Glencross, A case study evaluation: Perceptually accurate textured surface models, in: Proceedings of the 6th Symposium on Applied Perception in Graphics and Visualization, APGV '09, ACM, New York, NY, USA, 2009, pp. 109–115. doi:10.1145/1620993.1621015.
URL http://doi.acm.org/10.1145/1620993.1621015

[163] G. J. Székely, M. L. Rizzo, N. K. Bakirov, Measuring and testing dependence by correlation of distances, The Annals of Statistics 35 (6) (2007) 2769–2794.
URL http://www.jstor.org/stable/25464608

[164] D. H. Brainard, S. Vision, The psychophysics toolbox, Spatial vision 10 (1997) 433–436.

[165] D. G. Pelli, The videotoolbox software for visual psychophysics: Transforming numbers into movies, Spatial vision 10 (4) (1997) 437–442.

[166] M. Kleiner, D. Brainard, D. Pelli, What's new in psychtoolbox-3?, Vol. 36, 2007, p. 1.

[167] U. I. for Creative Technologies, High-resolution light probe image gallery (2017).
URL `http://gl.ict.usc.edu/Data/HighResProbes/`

[168] CIE, 116-1995 industrial colour-difference evaluation, Tech. rep.

[169] A. Sztrajman, J. Krivanek, A. Wilkie, T. Weyrich, Image-based remapping of material appearance, in: R. Klein, H. Rushmeier (Eds.), Proc. 5th Workshop on Material Appearance Modeling, MAM '17, The Eurographics Association, Aire-la-Ville, Switzerland, Switzerland, 2017.

[170] L. Kettner, M. Raab, D. Seibert, J. Jordan, A. Keller, The Material Definition Language, in: R. Klein, H. Rushmeier (Eds.), Workshop on Material Appearance Modeling, The Eurographics Association, 2015. doi:10.2312/mam.20151195.

[171] S. Bell, P. Upchurch, N. Snavely, K. Bala, Material recognition in the wild with the materials in context database, in: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2015.

[172] S. Onstott, Adobe Photoshop CS6 Essentials, 1st Edition, SYBEX Inc., Alameda, CA, USA, 2012.

[173] D. Shapira, S. Avidan, Y. Hel-Or, Multiple histogram matching.

[174] Y. Liu, M. Cohen, M. Uyttendaele, S. Rusinkiewicz, Autostyle: Automatic style transfer from image collections to users images, Computer Graphics Forum (Proc. Eurographics Symposium on Rendering 33 (4).

[175] S. Bell, P. Upchurch, N. Snavely, K. Bala, Material recognition in the wild with the materials in context database, in: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2015.

[176] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, A. Rabinovich, Going deeper with convolutions, in: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2015.

[177] B. E. Rogowitz, T. Frese, J. R. Smith, C. A. Bouman, E. B. Kalin, Perceptual image similarity experiments., in: Human Vision and Electronic Imaging, Vol. 3, 1998, p. 1.

[178] M. Toscani, M. Valsecchi, K. R. Gegenfurtner, Lightness perception for matte and glossy complex shapes, Vision Research 131 (Supplement C) (2017) 82 – 95. doi:https://doi.org/10.1016/j.visres.2016.12.004.

[179] S. A. Cholewiak, R. W. Fleming, Towards a unified explanation of shape from shading and texture, Journal of Vision 13 (9) (2013) 258–258.

[180] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, A. Rabinovich, Going deeper with convolutions, in: 2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2015, pp. 1–9. doi:10.1109/CVPR.2015.7298594.